arXiv:1703.09848v1 [cs.IT] 29 Mar 2017

Painless Breakups – Efficient Demixing of Low Rank Matrices

Thomas Strohmer and Ke Wei*
Department of Mathematics
University of California at Davis
Davis, CA 95616
{strohmer,kewei}@math.ucdavis.edu

March 30, 2017

Abstract

Assume we are given a sum of linear measurements of s different rank-r matrices of the form $\boldsymbol{y} = \sum_{k=1}^{s} \mathcal{A}_k(\boldsymbol{X}_k)$. When and under which conditions is it possible to extract (demix) the individual matrices $\boldsymbol{X}_k$ from the single measurement vector $\boldsymbol{y}$? And can we do the demixing numerically efficiently? We present two computationally efficient algorithms based on hard thresholding to solve this low rank demixing problem. We prove that under suitable conditions these algorithms are guaranteed to converge to the correct solution at a linear rate. We discuss applications in connection with quantum tomography and the Internet-of-Things. Numerical simulations demonstrate the empirical performance of the proposed algorithms.

1 Introduction

Demixing problems appear in a wide range of areas, including audio source separation [16], image processing [3], communications engineering [25], and astronomy [20]. A vast amount of recent literature focuses on the demixing of signals with very different properties. For instance one may be interested in the demixing of a signal that is the sum of spikes and sines [17] or the separation of a matrix that is the sum of a sparse matrix and a low rank matrix [4]. This paper focuses on the demixing problem where the signals to be separated all have *similar* properties, i.e., they are all low rank matrices. Assume we are given a sum of linear measurements

$$\boldsymbol{y} = \sum_{k=1}^{s} \mathcal{A}_k(\boldsymbol{X}_k), \tag{1}$$

*The authors acknowledge support from the NSF via grants dtra-dms 1322393 and DMS 1620455.

where $\{\mathcal{A}_k\}_{k=1}^s$ is a set of linear operators from $n \times n$ matrices to m-dimensional vectors and the constituents $\{\boldsymbol{X}_k\}_{k=1}^s$ are unknown rank-r matrices. Our goal is to extract each constituent matrix $\boldsymbol{X}_k$ from the single observation vector $\boldsymbol{y}$. We face different challenges in this homogeneous scenario, and one way to make the separation possible is by ensuring that the sensing operators $\mathcal{A}_k$ are sufficiently different from each other. The problem of demixing a sum of low rank matrices arises in quantum tomography [8], dictionary learning [12] and wireless communications [29].

In quantum tomography one tries to reconstruct an unknown quantum state from experimental data [13]. The state of a quantum system in quantum mechanics is often described by a density matrix, which is a positive semi-definite Hermitian matrice with unit trace. Many density matrices of interest are (or at least approximately) low rank—for instance pure states can be represented by rank-one matrices. Hence, in quantum tomography the goal is to reconstruct low rank matrices from a set of linear measurements, and ideally by using as few measurements as possible. In [8], a specific measurement protocol for quantum tomography is described in which states get mixed together. In this case one has to reconstruct *and* demix quantum states (i.e., low rank matrices) from a set of linear measurements.

Demixing problems of the form (1) are also expected to arise in the future Internet-of-Things (IoT). The IoT will connect billions of wireless devices, which is far more than the current wireless systems can technically and economically accommodate. One of the many challenges in the design of the IoT will be its ability to manage the massive number of *sporadic traffic* generating devices which are inactive most of the time, but regularly access the network for minor updates with no human interaction [29]. It is common understanding among communication engineers that this traffic cannot be handled within the current random access procedures. Dimensioning the channel access according to classical information and communication theory results in a severe waste of resources which does not scale towards the requirements of the IoT. This means among others that the overhead caused by the exchange of certain types of information between transmitter and receiver, such as channel estimation, assignment of data slots, etc, has to be avoided as much as possible. Without explicit channel estimation the receiver needs to blindly deconvolve (to undo the effect of the channel) and demix the signals arriving from many different devices. We will describe in Section 3 how this blind deconvolution-demixing problem can be phrased as a demixing problem of rank-one matrices.

1.1 State of the Art

The low rank matrix demixing problem in (1) is a generalization of the well-known low rank matrix recovery problem [19]. The task in low rank matrix recovery is to reconstruct a single low rank matrix $\boldsymbol{X}$ from a few linear measurements $\boldsymbol{y} = \mathcal{A}(\boldsymbol{X})$; that is, we have $s = 1$ in (1). This problem finds applications in a wide range of disciplines, including quantum tomography [13], and image processing [19, 30]. Many different algorithms have been proposed for its solution [30], including convex optimization based methods [6] and non-convex methods such as thresholding algorithms [2, 22]. The thresholding-based algorithms proposed in this paper can be viewed as natural generalizations of the latter algorithms.

There is a plethora of literature of various demixing problems, such as demixing of a low rank matrix and a sparse matrix, see for example [4]. However, most of these problems differ considerably from the one studied in this paper, since they focus on the separation of two signals with complementary properties. Therefore we will not discuss them here in any further detail.

Two of the first papers that consider the demixing of a sum of low rank matrices are [28] and [18]. In [18], the authors consider a general demixing framework of the form

$$\boldsymbol{y} = \boldsymbol{A}\left(\sum_{k=1}^{s} \boldsymbol{U}_k \boldsymbol{X}_k + \boldsymbol{w}_k\right), \tag{2}$$

where $\boldsymbol{A}$ is a right-invertible matrix, the matrices $\{\boldsymbol{U}_k\}_{k=1}^s$ are random unitary matrices, the matrices $\{\boldsymbol{X}_k\}_{k=1}^s$ are the signals of interest and the vectors $\{\boldsymbol{w}_k\}_{k=1}^s$ are noise. In the framework by McCoy and Tropp the signals $\boldsymbol{X}_k$ are assumed to be highly structured, which includes the scenario where all the $\boldsymbol{X}_k$ are rank-r matrices. Thus, in that case the setup (2) is a special instance of (1). The focus of [18] is on theoretical performance bounds of convex approaches for solving (2). In a nutshell the authors derive compelling quantitative phase transition bounds which show that demixing can succeed if and only if the dimension of the observation exceeds the total degrees of freedom present in the signals.

A special case of (1) where the constituents $\{\boldsymbol{X}_k\}_{k=1}^s$ are rank-one matrices has been analyzed from a theoretical and a numerical viewpoint in [14]. There, the authors investigate nuclear norm minimization (NNM) for the demixing problem under structured measurements of practical interest, and show that $m \gtrsim s^2 n$ number of measurements are sufficient for NNM to reliably extract each constituent $\boldsymbol{X}_k$ from the single measurement vector $\boldsymbol{y}$. It is also worth noting that an improvement of the theoretical analysis in [14] has been announced in [21].

A limitation of the numerical methods proposed in [14] and [18] is that the resulting semidefinite program is computationally rather expensive to solve for medium-size and large-size problems. Some applications require numerical efficiency, in which case a different numerical approach is needed. The goal of this paper is to develop numerically efficient algorithms that can solve the nonlinear low rank matrix demixing problem without resorting to convex optimization, and meanwhile to provide competitive theoretical recovery guarantees for the proposed algorithms. Closest to the goal of this paper is arguably a very recent paper by Ling and one of the authors [15]. There, the authors consider the same setup as in [14], and propose a non-convex regularized gradient-descent based method. The differences to this paper are that (i) [15] is specialized to the joint blind deconvolution-demixing setting and is not designed for (1), where the unknown matrices $\boldsymbol{X}_k$ are general rank-r matrices and the linear operators $\mathcal{A}_k$ are more general sensing matrices; (ii) while both algorithms fall in the realm of non-convex optimization, [15] uses a gradient descent method and this paper uses thresholding-based methods; (iii) the theoretical analysis in this paper does not apply to the joint blind deconvolution-demixing setting in [15].

1.2 Outline and Notation

The remainder of the paper is organized as follows. The numerical algorithms and their theoretical guarantees for the demixing problem are presented in Section 2. We also introduce an Amalgam form of restricted isometry property in Section 2 around which our theoretical analysis revolves. In Section 3 we test our algorithms on a variety of numerical examples. The proofs of the theoretical results are presented in Section 4. We conclude this paper with some potential future directions in Section 5.

Throughout the paper we use the following notational conventions. We denote vectors by bold lowercase letters and matrices by bold uppercase letters. In particular, we fix $\{\boldsymbol{X}_k\}_{k=1}^s$ as the target matrices and $\{\boldsymbol{X}_{k,l}\}_{k=1}^s$ as the iterates of the algorithms. For a matrix $\boldsymbol{Z}$, we use $\|\boldsymbol{Z}\|$ and $\|\boldsymbol{Z}\|_F$ to denote its spectral norm and Frobenius norm, respectively. For both vectors and matrices, $\boldsymbol{z}^T$ and $\boldsymbol{Z}^T$ denote their transpose while $\boldsymbol{z}^*$ and $\boldsymbol{Z}^*$ denote their conjugate transpose. The inner product of two matrices $\boldsymbol{Z}_1$ and $\boldsymbol{Z}_2$ is defined as $\langle \boldsymbol{Z}_1, \boldsymbol{Z}_2 \rangle = \operatorname{trace}(\boldsymbol{Z}_1^* \boldsymbol{Z}_2)$.

2 Amalgam-RIP and Algorithms

2.1 Problem Setup and Amalgam-RIP

As outlined in the introduction, we want to solve the following demixing problem:

$$\text{Find all rank-}r\text{ matrices } \boldsymbol{X}_k, \qquad \text{given } \boldsymbol{y} = \sum_{k=1}^{s} \mathcal{A}_k(\boldsymbol{X}_k), \tag{3}$$

where $\{\mathcal{A}_k\}_{k=1}^s$ is a set of linear operators mapping $n \times n$ matrices to m-dimensional vectors. Let $\{\boldsymbol{A}_{k,p} \mid 1 \le k \le s, 1 \le p \le m\}$ be a set of measurement matrices. We can write $\mathcal{A}_k(\boldsymbol{Z})$ explicitly as

$$\mathcal{A}_k(\boldsymbol{Z}) = \frac{1}{\sqrt{m}} \begin{bmatrix} \langle \boldsymbol{A}_{k,1}, \boldsymbol{Z} \rangle \\ \vdots \\ \langle \boldsymbol{A}_{k,m}, \boldsymbol{Z} \rangle \end{bmatrix}. \tag{4}$$

For conciseness, we only present our algorithms and theoretical results for $n \times n$ real matrices, but it is straightforward to modify them for $n_1 \times n_2$ complex matrices.

We will propose two iterative hard thresholding algorithms for low rank matrix demixing. A question of central importance is how many measurements are needed so that the algorithms can successfully extract all the constituents $\{\boldsymbol{X}_k\}_{k=1}^s$ from $\boldsymbol{y}$. Since each $\boldsymbol{X}_k$ is determined by $r(2n-r)$ parameters [5], we need at least $m \ge sr(2n-r)$ many measurements for the problem to be well-posed. One of our goals is to show that the proposed algorithms can succeed already if the number of measurements is close to this information-theoretic minimum. As is common for hard thresholding algorithms in compressed sensing and low rank matrix recovery, the convergence analysis here will rely on some form of restricted

isometry property (RIP). The notion most fitting to the demixing setting requires the RIP in an amalgamated form.

Definition 2.1 (Amalgam-RIP)**.** *The set of linear operators $\{\mathcal{A}_k\}_{k=1}^s$ satisfy the Amalgam-RIP (ARIP) with the parameter δ_r if*

$$(1-\delta_r)\sum_{k=1}^{s}\|\boldsymbol{Z}_k\|_F^2 \le \left\|\sum_{k=1}^{s}\mathcal{A}_k(\boldsymbol{Z}_k)\right\|^2 \le (1+\delta_r)\sum_{k=1}^{s}\|\boldsymbol{Z}_k\|_F^2 \tag{5}$$

holds for all the matrices $\boldsymbol{Z}_k,\ k=1,\cdots,s$ of rank at most r.

Fix an index k and take $\boldsymbol{Z}_{k'}=0$ for all $k'\neq k$. The ARIP implies that

$$(1-\delta_r)\|\boldsymbol{Z}_k\|_F^2 \le \|\mathcal{A}_k(\boldsymbol{Z}_k)\|^2 \le (1+\delta_r)\|\boldsymbol{Z}_k\|_F^2$$

hold for all matrices $\boldsymbol{Z}_k$ of rank at most r, which is indeed the RIP introduced in [19] for low rank matrix recovery. When the measurements matrices in (4) have i.i.d $\mathcal{N}(0,1)$ entries, we can show that $\{\mathcal{A}_k\}_{k=1}^s$ satisfy the ARIP with overwhelmingly high probability provided the number of measurements is proportional to the number of degrees of freedom within the constituents $\{\boldsymbol{X}_k\}_{k=1}^s$.

Theorem 2.1. *If $\{\boldsymbol{A}_{k,p} \mid 1\le k\le s, 1\le p\le m\}$ is a set of standard Gaussian matrices, then $\{\mathcal{A}_k\}_{k=1}^s$ satisfy the ARIP with the parameter δ with probability at least $1-2\exp(-m\delta^2/64)$ provided*

$$m \ge C\delta^{-2}(2n+1)rs\log(1/\delta),$$

where $C>0$ is an absolute numerical constant.

2.2 Iterative Hard Thresholding Algorithms

In this subsection, we present two iterative hard thresholding algorithms for the low rank demixing problem. The algorithms are developed by targeting the following rank constraint problem directly

$$\min_{\{\boldsymbol{Z}_k\}_{k=1}^s}\left\|\boldsymbol{y}-\sum_{k=1}^{s}\mathcal{A}_k(\boldsymbol{Z}_k)\right\|^2 \quad \text{subject to} \quad \operatorname{rank}(\boldsymbol{Z}_k)\le r.$$

The first algorithm, simply referred to as iterative hard thresholding (IHT), is presented in Algorithm 1. In each iteration, the algorithm first computes the residual and then updates each constituent separately by projected gradient descent. The search direction with respect to each constituent is computed as the negative gradient descent direction of the objective function when the other constituents are fixed. The hard thresholding operator $\mathcal{H}_r$ in the algorithm returns the best rank r approximation of a matrix. The stepsize within the line search can be either fixed or computed adaptively. In Algorithm 1, we compute it as the

steepest descent stepsize along the projected gradient descent direction $\mathcal{P}_{T_{k,l}}(\boldsymbol{G}_{k,l})$, where $T_{k,l}$ is the tangent space of the rank r matrix manifold at the current estimate $\boldsymbol{X}_{k,l}$. Let $\boldsymbol{X}_{k,l} = \boldsymbol{U}_{k,l}\boldsymbol{\Sigma}_{k,l}\boldsymbol{V}_{k,l}^*$ be the singular value decomposition of $\boldsymbol{X}_{k,l}$. The tangent space $T_{k,l}$ consists of matrices which share the same column or row subspaces with $\boldsymbol{X}_{k,l}$ [1],

$$T_{k,l} = \{\boldsymbol{U}_{k,l}\boldsymbol{Z}_1^* + \boldsymbol{Z}_2\boldsymbol{V}_{k,l}^* \mid \boldsymbol{Z}_1 \in \mathbb{R}^{n\times r},\ \boldsymbol{Z}_2 \in \mathbb{R}^{n\times r}\}. \tag{6}$$

Note that the inner loop within the algorithm is fully parallel since we update each constituent separately.

Algorithm 1 Iterative Hard Thresholding for Low Rank Demixing

Initialization: $\boldsymbol{X}_{k,0} = \boldsymbol{U}_{k,0}\boldsymbol{\Sigma}_{k,0}\boldsymbol{V}_{k,0}^* = \mathcal{H}_r(\mathcal{A}_k^*(\boldsymbol{y}))$
for $l = 0, 1, \cdots$ **do**
  $\boldsymbol{r}_l = \boldsymbol{y} - \sum_{k=1}^s \mathcal{A}_k(\boldsymbol{X}_{k,l})$
  for $k = 1, \cdots, s$ **do** {*Fully Parallel*}
    1. $\boldsymbol{G}_{k,l} = \mathcal{A}_k^*(\boldsymbol{r}_l)$
    2. $\alpha_{k,l} = \frac{\|\mathcal{P}_{T_{k,l}}(\boldsymbol{G}_{k,l})\|_F^2}{\|\mathcal{A}_k\mathcal{P}_{T_{k,l}}(\boldsymbol{G}_{k,l})\|_2^2}$
    3. $\boldsymbol{X}_{k,l+1} = \mathcal{H}_r(\boldsymbol{X}_{k,l} + \alpha_{k,l}\boldsymbol{G}_{k,l})$
  end for
end for

Theorem 2.2. *Assume the linear operators $\{\mathcal{A}_k\}_{k=1}^s$ satisfy the ARIP with the parameter δ_{3r}. Define*

$$\gamma_1 = \frac{4\delta_{3r}}{1-\delta_{3r}}.$$

If $\gamma_1 < 1$, then the iterates of IHT satisfy

$$\left\| \begin{bmatrix} \boldsymbol{X}_{1,l+1} - \boldsymbol{X}_1 \\ \vdots \\ \boldsymbol{X}_{s,l+1} - \boldsymbol{X}_s \end{bmatrix} \right\|_F \leq \gamma_1 \left\| \begin{bmatrix} \boldsymbol{X}_{1,l} - \boldsymbol{X}_1 \\ \vdots \\ \boldsymbol{X}_{s,l} - \boldsymbol{X}_s \end{bmatrix} \right\|_F.$$

It follows immediately that the iterates of IHT converge linearly to the underlying constituents provided $\delta_{3r} < 1/5$.

If $\{\mathcal{A}_k\}_{k=1}^s$ consist of standard Gaussian measurement matrices, then Theorem 2.2 together with Theorem 2.1 implies that the number of necessary measurements for IHT to achieve successful recovery is $O(snr)$ which is optimal. In addition, we can establish the robustness of IHT against additive noise by considering the model $\boldsymbol{y} = \sum_{k=1}^s \mathcal{A}_k(\boldsymbol{X}_k) + \boldsymbol{e}$, where $\boldsymbol{e}$ represents a noise term.

Theorem 2.3. *Assume* $\boldsymbol{y} = \sum_{k=1}^{s} \mathcal{A}_k(\boldsymbol{X}_k) + \boldsymbol{e}$ *and the linear operators* $\{\mathcal{A}_k\}_{k=1}^{s}$ *satisfy the ARIP with parameter* δ_{3r}*. If* $\gamma_1 < 1$*, then the iterates of IHT satisfy*

$$\left\| \begin{bmatrix} \boldsymbol{X}_{1,l} - \boldsymbol{X}_1 \\ \vdots \\ \boldsymbol{X}_{s,l} - \boldsymbol{X}_s \end{bmatrix} \right\|_F \leq \gamma_1^l \left\| \begin{bmatrix} \boldsymbol{X}_{1,0} - \boldsymbol{X}_1 \\ \vdots \\ \boldsymbol{X}_{s,0} - \boldsymbol{X}_s \end{bmatrix} \right\|_F + \frac{\xi}{1-\gamma_1} \|\boldsymbol{e}\|,$$

where

$$\xi = \frac{2\sqrt{1+\delta_{3r}}}{1-\delta_{2r}}.$$

The application of the hard thresholding operator in Algorithm 1 requires the singular value decomposition (SVD) of an $n \times n$ matrix in each iteration, which is computationally expensive for unstructured matrices. Inspired by the work of Riemannian optimization for low rank matrix reconstruction in [23, 27, 26], we propose to accelerate IHT by updating each constituent along a projected gradient descent direction, see fast iterative hard thresholding (FIHT) described in Algorithm 2. The key difference between Algorithms 1 and 2 lies in Step 3. Instead of updating $\boldsymbol{X}_{k,l}$ along the gradient descent direction $\boldsymbol{G}_{k,l}$ as in IHT, FIHT updates $\boldsymbol{X}_{k,l}$ along the projected gradient descent direction $\mathcal{P}_{T_{k,l}}(\boldsymbol{G}_{k,l})$, where $T_{k,l}$ is the tangent space defined in (6), and then followed by the hard thresholding operation.

Algorithm 2 Fast Iterative Hard Thresholding for Low Rank Demixing

Initialization: $\boldsymbol{X}_{k,0} = \boldsymbol{U}_{k,0}\boldsymbol{\Sigma}_{k,0}\boldsymbol{V}_{k,0}^* = \mathcal{H}_r(\mathcal{A}_k^*(\boldsymbol{y}))$
for $l = 0, 1, \cdots$ **do**
  $\boldsymbol{r}_l = \boldsymbol{y} - \sum_{k=1}^{s} \mathcal{A}_k(\boldsymbol{X}_{k,l})$
  for $k = 1, \cdots, s$ **do** {*Fully Parallel*}
    1. $\boldsymbol{G}_{k,l} = \mathcal{A}_k^*(\boldsymbol{r}_l)$
    2. $\alpha_{k,l} = \frac{\|\mathcal{P}_{T_{k,l}}(\boldsymbol{G}_{k,l})\|_F^2}{\|\mathcal{A}_k\mathcal{P}_{T_{k,l}}(\boldsymbol{G}_{k,l})\|_2^2}$
    3. $\boldsymbol{X}_{k,l+1} = \mathcal{H}_r(\boldsymbol{X}_{k,l} + \alpha_{k,l}\mathcal{P}_{T_{k,l}}(\boldsymbol{G}_{k,l}))$
  end for
end for

Let $\boldsymbol{W}_{k,l} = \boldsymbol{X}_{k,l} + \alpha_{k,l}\mathcal{P}_{T_{k,l}}(\boldsymbol{G}_{k,l})$. One can easily observe that $\boldsymbol{W}_{k,l} \in T_{k,l}$ and $\operatorname{rank}(\boldsymbol{W}_{k,l}) \leq 2r$. Moreover, $\boldsymbol{W}_{k,l}$ obeys the following decomposition [27]:

$$\boldsymbol{W}_{k,l} = \begin{bmatrix} \boldsymbol{U}_{k,l} & \boldsymbol{Q}_1 \end{bmatrix} \boldsymbol{M}_{k,l} \begin{bmatrix} \boldsymbol{V}_{k,l}^* \\ \boldsymbol{Q}_2^* \end{bmatrix}$$

where $\boldsymbol{Q}_1$ and $\boldsymbol{Q}_2$ are two $n \times r$ orthonormal matrices such that $\boldsymbol{Q}_1 \perp \boldsymbol{U}_{k,l}$ and $\boldsymbol{Q}_2 \perp \boldsymbol{V}_{k,l}$, and $\boldsymbol{M}_{k,l}$ is a $2r \times 2r$ matrix of the form

$$\boldsymbol{M}_{k,l} = \begin{bmatrix} \boldsymbol{\Sigma}_{k,l} + \boldsymbol{\xi}_{k,l} & \boldsymbol{R}_1^* \\ \boldsymbol{R}_2^* & 0 \end{bmatrix}.$$

Consequently, the SVD of $\boldsymbol{W}_{k,l}$ can be obtained from the SVD of the $2r \times 2r$ matrix $\boldsymbol{M}_{k,l}$. Therefore, the introduction of the extra projection $\mathcal{P}_{T_{k,l}}$ for the search direction can reduce the computational complexity of the partial SVD of $\boldsymbol{W}_{k,l}$ from $O(n^2 r)$ flops to $O(nr^2 + r^3)$ flops. The recovery guarantee of FIHT can also be established in terms of the ARIP.

Theorem 2.4. *Assume the linear operators $\{\mathcal{A}_k\}_{k=1}^s$ satisfy the ARIP with the parameter δ_{3r}. Define*

$$\gamma_2 = 2\left(\frac{2\delta_{2r}}{1-\delta_{2r}} + \frac{\delta_{3r}}{1-\delta_{2r}} + 2\delta_{3r}\sqrt{rs}\frac{\sigma_{\max}}{\sigma_{\min}}\right),$$

where $\sigma_{\max} := \max_k \sigma_{\max}(\boldsymbol{X}_k)$ and $\sigma_{\min} := \max_k \sigma_{\min}(\boldsymbol{X}_k)$. If $\gamma_2 < 1$, then the iterates of Algorithm 2 satisfy

$$\left\| \begin{bmatrix} \boldsymbol{X}_{1,l+1} - \boldsymbol{X}_1 \\ \vdots \\ \boldsymbol{X}_{s,l+1} - \boldsymbol{X}_s \end{bmatrix} \right\|_F \leq \gamma_2 \left\| \begin{bmatrix} \boldsymbol{X}_{1,l} - \boldsymbol{X}_1 \\ \vdots \\ \boldsymbol{X}_{s,l} - \boldsymbol{X}_s \end{bmatrix} \right\|_F.$$

It follows that the iterates of FIHT converge linearly to the underlying constituents provided

$$\delta_{3r} \lesssim \frac{\sigma_{\min}}{\sigma_{\max}} \frac{1}{\sqrt{rs}}.$$

Assume $\{\mathcal{A}_k\}_{k=1}^s$ consist of standard Gaussian matrices. In contrast to the $O(srn)$ necessary measurements for IHT to achieve successful recovery, Theorem 2.4 implies that FIHT requires $O((\sigma_{\max}/\sigma_{\min})^2 s^2 nr^2)$ measurements which is not optimal. However, numerical simulations in Section 3 suggest that FIHT needs fewer measurements than IHT to be able to successfully reconstruct s rank-r matrices.

3 Numerical Simulations and Applications

We evaluate the empirical performance of IHT and FIHT on simulated problems as well as application examples from quantum states demxing and the Internet of Things. In the implementation IHT and FIHT are terminated when a maximum of 500 iterations is met or the relative residual is small,

$$\frac{\|\boldsymbol{y} - \sum_{k=1}^s \mathcal{A}_k(\boldsymbol{X}_{k,l})\|}{\|\boldsymbol{y}\|} \leq 10^{-4}.$$

They are considered to have successfully recovered a set of matrices $\{\boldsymbol{X}_k\}_{k=1}^s$ if the returned approximations $\{\boldsymbol{X}_{k,l}\}_{k=1}^s$ satisfy

$$\frac{\sqrt{\sum_{k=1}^s \|\boldsymbol{X}_{k,l} - \boldsymbol{X}_k\|_F^2}}{\sqrt{\sum_{k=1}^s \|\boldsymbol{X}_k\|_F^2}} \leq 10^{-2}.$$

All the random tests are repeated *ten* times in this section.

3.1 Phase Transitions under Gaussian Measurements

The ARIP-based theoretical recovery guarantees in Theorems 2.2 and 2.4 are worst-case analysis which are uniform over all rank r matrices. However, these conditions are highly pessimistic when compared with average-case empirical observations. In this subsection we investigate the empirical recovery performance of IHT and FIHT under Gaussian measurements. The tests are conducted on $n \times n$ rank-r matrices with $n = 50$ and $r = 5$, and the number of constituents s varies from 1 to 7. For a fixed value of constituents, a number of equispaced values of m are tested.

We first test IHT and FIHT on well-conditioned matrices which are formed via $\boldsymbol{X}_k = \boldsymbol{L}_k \boldsymbol{R}_k^*$, where $\boldsymbol{L}_k \in \mathbb{R}^{n\times r}$, $\boldsymbol{R}_k \in \mathbb{R}^{n\times r}$, and their entries are drawn i.i.d. from the standard normal distribution. The phase transitions of IHT and FIHT on well-conditioned matrices are presented in Figures 1a and 1b, respectively. Figure 1a shows a linear correlation between the number of measurements m and the number of constituents s for the successful recovery of IHT. Though the recovery guarantee of FIHT in Theorem 2.4 is worse than that of IHT in Theorem 2.2, Figure 1b shows that FIHT requires fewer measurements than IHT to achieve successful recovery with high probability. In particular, $m \approx 3.3sr(2n - r)$ is sufficient for FIHT to successfully reconstruct a set of s low rank matrices for all the tested values of s.

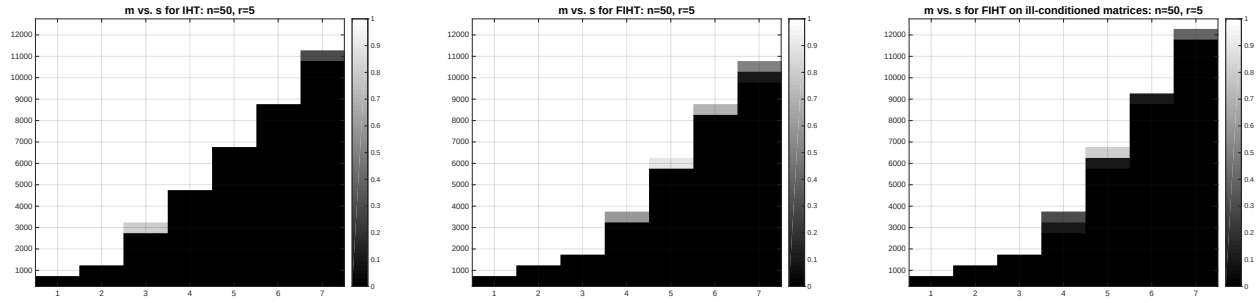

(a) (b) (c)

Figure 1: Empirical phase transitions of IHT and FIHT under Gaussian measurements.

To explore how the condition number of the test matrices impact the recovery of FIHT, we further test FIHT on ill-conditioned matrices. The matrices are computed as $\boldsymbol{X}_k = \boldsymbol{U}_k \boldsymbol{\Sigma} \boldsymbol{V}_k^*$, where $\boldsymbol{U}_k$ and $\boldsymbol{V}_k$ are two 50 by 5 random orthonormal matrices, and $\boldsymbol{\Sigma}$ is a 5 by 5 diagonal matrix with the diagonal entries being 5 uniformly spaced values from 1 to 1000. So the condition number of the test matrices is 1000. The phase transition plot of FIHT on ill-conditioned matrices is presented in Figure 1c. The figure shows that FIHT requires more measurements to reconstruct the test ill-conditioned matrices than to reconstruct the well-conditioned matrices when $s \geq 5$. However, FIHT is still able to reconstruct all the test ill-conditioned matrices when $m \approx 3.8sr(2n - r)$. Hence, there is no strong evidence from Figure 1c to support that the number of measurements for the successful recovery of FIHT relies on the square of the condition number of the target matrices, which may suggest the possibility of improving the result in Theorem 2.4.

In the reminder of this section, we focus on the empirical performance of FIHT for its superiority and flexibility. Though we have only established the recovery guarantee of FIHT for real matrices under the Gaussian measurements, the algorithm is equally effective for complex matrices and other types of measurements. In the next two subsections, we study the applications of FIHT in quantum states demixing and the Internet of Things.

3.2 An Application in Quantum States Demixing

As stated in the introduction, the state of a quantum system can be represented by a low rank, positive semidefinite density matrix of unit trace. Next, we present a stylized application of FIHT to the demixing of quantum states. When the target constituents are n by n rank-r positive semidefinite matrices of unit trace, FIHT in Algorithm 2 can be modified accordingly to preserve the matrix structures, see Algorithm 3 for complete description.

Algorithm 3 Fast Iterative Hard Thresholding for Quantum States Demixing

Initialization: $\boldsymbol{X}_{k,0} = \boldsymbol{U}_{k,0}\boldsymbol{\Lambda}_{k,0}\boldsymbol{U}_{k,0}^* = \mathcal{P}_\Delta\left(\mathcal{P}_\Pi(\mathcal{A}_k^*(\boldsymbol{y}))\right)$
for $l = 0, 1, \cdots$ **do**
  $\boldsymbol{r}_l = \boldsymbol{y} - \sum_{k=1}^s \mathcal{A}_k(\boldsymbol{X}_{k,l})$
  for $k = 1, \cdots, s$ **do** $\{$*Fully Parallel*$\}$
    1. $\boldsymbol{G}_{k,l} = \mathcal{A}_k^*(\boldsymbol{r}_l)$
    2. $\alpha_{k,l} = \frac{\|\mathcal{P}_{T_{k,l}}(\boldsymbol{G}_{k,l})\|_F^2}{\|\mathcal{A}_k\mathcal{P}_{T_{k,l}}(\boldsymbol{G}_{k,l})\|_2^2}$
    3. $\boldsymbol{X}_{k,l+1} = \mathcal{P}_\Delta\left(\mathcal{P}_\Pi(\boldsymbol{X}_{k,l} + \alpha_{k,l}\mathcal{P}_{T_{k,l}}(\boldsymbol{G}_{k,l}))\right)$
  end for
end for

Let $\boldsymbol{X}_{k,l}$ be the current rank-r estimate which is Hermitian, positive semidefinite, and so admits an eigenvalue decomposition (equivalent to its SVD) $\boldsymbol{X}_{k,l} = \boldsymbol{U}_{k,l}\boldsymbol{\Lambda}_{k,l}\boldsymbol{U}_{k,l}^*$ with non-negative eigenvalues. There are two key differences between Algorithms 2 and 3. Firstly, the tangent space can be defined as

$$T_{k,l} = \{\boldsymbol{U}_{k,l}\boldsymbol{Z}^* + \boldsymbol{Z}\boldsymbol{U}_{k,l}^* \mid \boldsymbol{Z} \in \mathbb{C}^{n\times r}\}.$$

so that all the matrices in $T_{k,l}$ are Hermitian. Therefore after $\boldsymbol{X}_{k,l}$ is updated along the projected gradient descent direction, $\boldsymbol{W}_{k,l} = \boldsymbol{X}_{k,l} + \alpha_{k,l}\mathcal{P}_{T_{k,l}}(\boldsymbol{G}_{k,l})$ remains Hermitian. In addition, the eigenvalue decomposition of $\boldsymbol{W}_{k,l}$ can also be computed using $O(nr^2)$ flops as in the computation of the SVD of the non-symmetric matrices in Algorithm 2, since $\boldsymbol{W}_{k,l}$ has the following decomposition:

$$\boldsymbol{W}_{k,l} = \begin{bmatrix}\boldsymbol{U}_{k,l} & \boldsymbol{Q}\end{bmatrix} \boldsymbol{M}_{k,l} \begin{bmatrix}\boldsymbol{U}_{k,l}^* \\ \boldsymbol{Q}^*\end{bmatrix},$$

where $\begin{bmatrix}\boldsymbol{U}_{k,l} & \boldsymbol{Q}\end{bmatrix}$ is an $n \times 2r$ orthonormal matrix and $\boldsymbol{M}_{k,l}$ is a $2r \times 2r$ Hermitian matrix of the form

$$\boldsymbol{M}_{k,l} = \begin{bmatrix}\boldsymbol{\Lambda}_{k,l} + \boldsymbol{\xi}_{k,l} & \boldsymbol{R}^* \\ \boldsymbol{R}^* & 0\end{bmatrix}.$$

Secondly, the hard thresholding operator $\mathcal{H}_r$ is replaced by the projection $\mathcal{P}_\Pi$, followed by another projection $\mathcal{P}_\Delta$, where $\mathcal{P}_\Pi$ projects $\boldsymbol{W}_{k,l}$ onto the set of low rank positive semidefinite matrices

$$\Pi = \{\boldsymbol{Z} \in \mathbb{C}^{n\times n} \mid \boldsymbol{Z} = \boldsymbol{Z}^*, \mathrm{rank}(\boldsymbol{Z}) \leq r, \ \text{ and } \boldsymbol{Z} \succeq 0\};$$

and $\mathcal{P}_S$ projects the eigenvalues of $\mathcal{P}_\Pi(\boldsymbol{W}_{k,l})$ onto the simplex

$$\Delta = \left\{ \boldsymbol{x} \in \mathbb{R}^r \mid \sum_{i=1}^{r} x_i = 1, \ x_i \geq 0 \right\}.$$

Notice that the inertia of a Hermitian matrix $\boldsymbol{H}$, denoted by inertia($\boldsymbol{H}$), is a triple of integer numbers (i_p, i_n, i_z), where i_p, i_n, i_z are the number of positive, negative, and zero eigenvalues of $\boldsymbol{H}$. From [9, Lemma A.15], we know that inertia($\boldsymbol{M}_{k,l}$) = inertia($\boldsymbol{N}^*\boldsymbol{M}_{k,l}\boldsymbol{N}$)+ $(t, t, r-t)$, where $t = \mathrm{rank}(\boldsymbol{R})$ and $\boldsymbol{N}$ is an $r\times(r-t)$ matrix whose columns forms a null space of $\boldsymbol{R}$. This implies that $\boldsymbol{M}_{k,l}$ has at most r positive eigenvalues, so does $\boldsymbol{W}_{k,l}$. Therefore by [11, Theorem 2.1], we can compute $\mathcal{P}_\Pi(\boldsymbol{W}_{k,l})$ by only keeping the positive eigenvalues and the corresponding eigenvectors of $\boldsymbol{M}_{k,l}$. To compute $\mathcal{P}_S(\cdot)$, we use an algorithm proposed in [7], in which the computational cost is dominated by sorting a vector of length r.

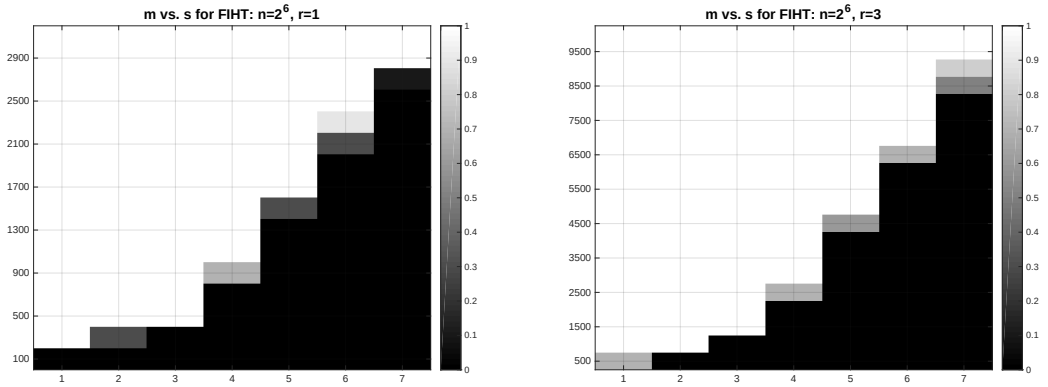

(a) (b)

Figure 2: Empirical phase transitions of FIHT for quantum states demixing.

We evaluate the recovery performance of FIHT for quantum states demixing under the Pauli measurements, where the measurement matrices in (4) are tensor products of the Pauli matrices. That is,

$$\boldsymbol{A}_{k,p} = 2^{-q/2} \otimes_{i=1}^{q} \sigma_a,$$

where

$$\sigma_1 = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}, \ \sigma_2 = \begin{bmatrix} 0 & -i \\ i & 0 \end{bmatrix}, \ \sigma_3 = \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}, \ \sigma_4 = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}.$$

The target matrices are computed as $\boldsymbol{X}_{k,l} = \frac{1}{r}\boldsymbol{U}_{k,l}\boldsymbol{U}_{k,l}^*$, where $\boldsymbol{U}_{k,l}$ are random $n \times r$ complex orthonormal matrices. We conduct the experiments with $q = 6$, $r = 1$ and 3, and a is selected uniformly at random. When $r = 1$, it corresponds to demixing the pure states in quantum mechanics. The phase transitions of FIHT under the Pauli measurements is presented in Figure 2, which shows that FIHT is equally effective for quantum states demixing and $m \approx 3.5sr(2n - r)$ is sufficient for successful demixing all the tested values of s from 1 to 7.

3.3 An Application from the Internet-of-Things

Demixing problems of the form (1) are also expected to play an important role in the future IoT. In this subsection we conduct some numerical simulations for a demixing problem related to the IoT. Before we proceed to the simulations we describe in more detail how such a demixing problem does arise in the IoT and how it is connected to the demixing of low rank matrices.

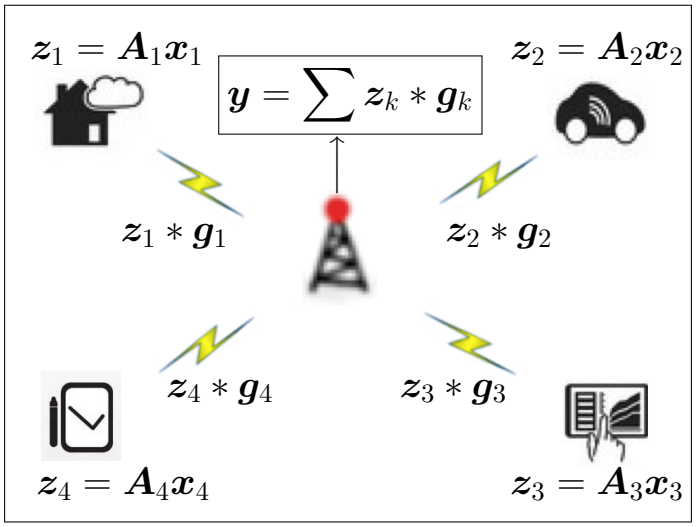

Figure 3: A random access wireless communication scenario in the Internet of Things: Each device communicates with a common basestation. No explicit channel estimation is done in order to reduce the scheduling overhead. The transmitted signals $\boldsymbol{z}_k$ are of the form $\boldsymbol{z}_k = \boldsymbol{C}_k\boldsymbol{x}_k$, where the $\boldsymbol{C}_k$ are encoding matrices and the $\boldsymbol{x}_k$ represent the data to be transmitted. Each signal $\boldsymbol{z}_k$ travels through a wireless communication channel, represented by $\boldsymbol{g}_k$, which acts as convolution. The base station measures the signal $\boldsymbol{y}$, which is the sum of all received signals; and it knows the encoding matrices $\boldsymbol{C}_k$ of the individual users, but it does neither know the $\boldsymbol{z}_k$'s (or the $\boldsymbol{x}_k$'s) nor the transmission channels $\boldsymbol{g}_k$. The goal is to demix the received signal $\boldsymbol{y}$ and recover the signals $\boldsymbol{x}_k$ of the individual users.

In mathematical terms we are dealing—in somewhat simplified form—with the following problem. We consider a scenario where many different users/devices communicate with a base station, see Figure 3 for an illustration of the described setup. Let $\boldsymbol{x}_k \in \mathbb{C}^{n_1}$ be the

data that are supposed to be transmitted by the k-th user. Before transmission we encode $\boldsymbol{x}_k$ by computing $\boldsymbol{z}_k = \boldsymbol{C}_k \boldsymbol{x}_k$, where $\boldsymbol{C}_k \in \mathbb{C}^{m \times n_1}$ is a fixed precoding matrix of the k-th user. The encoding matrice $\boldsymbol{C}_k$ differs from user to user but is known to the receiver. Let $\boldsymbol{g}_k$ be the associated impulse response of the communication channel between the k-th user and the base station. Since we assume that the channel does not change during the transmission of the signal $\boldsymbol{z}_k$, $\boldsymbol{g}_k$ acts as convolution operator, i.e., the signal arriving at the base station becomes $\boldsymbol{z}_k * \boldsymbol{g}_k$ (ignoring additive noise for simplicity). However, the channels $\boldsymbol{g}_k$ are *not known* to the base station (or to the transmitter), since this is the overhead information that we want to avoid transmitting, as it may change from transmission to transmission. The signal recorded at the base station is given by a mixture of all transmitted signals convolved with their respective channel impulse responses, i.e.,

$$\boldsymbol{y} = \sum_k \boldsymbol{z}_k * \boldsymbol{g}_k. \tag{7}$$

The goal is to demix the received signal $\boldsymbol{y}$ and recover the signals $\boldsymbol{x}_k$ of the individual users.

But how is (7) related to our setup in (1)? We first note that in all wireless communication scenarios of relevance the impulse response $\boldsymbol{g}_k$ is compactly supported. In engineering terms the size of the support of $\boldsymbol{g}_k$, denoted here by n_2, is called its maximum delay spread. In practice, the length of the transmitted signals m is typically (much) larger than the maximum delay spread n_2. Thus, $\boldsymbol{g}_k$ is a m-dimensional vector which is formed by the channel impulse responses being padded with zeroes. We can write $\boldsymbol{g}_k$ as $\boldsymbol{g}_k = [\boldsymbol{h}_k^T, 0, \dots, 0]^T$, where $\boldsymbol{h}_k$ is the non-zero padded impulse response of length n_2. Let $\boldsymbol{F}$ denote the $m \times m$ unitary Discrete Fourier Transform matrix and let $\boldsymbol{B}$ be the $m \times n_2$ matrix consisting of the first n_2 columns of $\boldsymbol{F}$. By the circular convolution theorem, we can express the convolution between $\boldsymbol{z}_k$ and $\boldsymbol{g}_k$ in the Fourier domain via

$$\frac{1}{\sqrt{m}} \hat{\boldsymbol{y}} = (\boldsymbol{F}\boldsymbol{C}_k \boldsymbol{x}_k) \odot \boldsymbol{B}\boldsymbol{h}_k, \tag{8}$$

where $\odot$ denotes the componentwise product. Let $\boldsymbol{c}_{k,p}^*$ represent the p-th row of $\boldsymbol{F}\boldsymbol{C}_k$ and let $\boldsymbol{b}_p^*$ represent the p-th row of $\boldsymbol{B}$. A simple calculation shows that [14]

$$[(\boldsymbol{F}\boldsymbol{C}_k \boldsymbol{x}_k) \odot \boldsymbol{B}\boldsymbol{h}_k]_p = (\boldsymbol{c}_{k,p}^* \boldsymbol{x}_k) \cdot (\boldsymbol{b}_p^* \boldsymbol{h}_k) = \left\langle \boldsymbol{c}_{k,p} \bar{\boldsymbol{b}}_p^*, \boldsymbol{x}_k \bar{\boldsymbol{h}}_k^* \right\rangle = \left\langle \boldsymbol{c}_{k,p} \bar{\boldsymbol{b}}_p^*, \boldsymbol{X}_k \right\rangle, \tag{9}$$

where $\boldsymbol{X}_k := \boldsymbol{x}_k \bar{\boldsymbol{h}}_k^*$ is an $n_1 \times n_2$ rank-one matrix containing the unknown signal and the unknown channel of the k-th user, while $\boldsymbol{c}_{k,p}$ and $\boldsymbol{b}_p^*$ are known vectors (since $\boldsymbol{C}_k$ and $\boldsymbol{B}$ are known). Hence, at the base station the observed signal $\boldsymbol{y} = \sum_{k=1}^s \boldsymbol{z}_k * \boldsymbol{g}_k$ can be expressed in the Fourier domain as $\frac{1}{\sqrt{m}} \hat{\boldsymbol{y}} = \sum_{k=1}^s \mathcal{A}_k(\boldsymbol{X}_k)$, which is of the same form as (1) (modulo replacing $\frac{1}{\sqrt{m}} \hat{\boldsymbol{y}}$ with $\boldsymbol{y}$). The measurement matrices in this scenario are given by $\boldsymbol{A}_{k,p} = \boldsymbol{c}_{k,p} \bar{\boldsymbol{b}}_p^*$.

We test FIHT (Algorithm 2) for the demixing problem described in (8) with $n_1 = 128$ and $n_2 = 32$, which are of interest in practice. As a result, the constituents are non-square matrices. Two types of encoding matrices are tested: a) the entries of $\boldsymbol{C}_k$ are standard

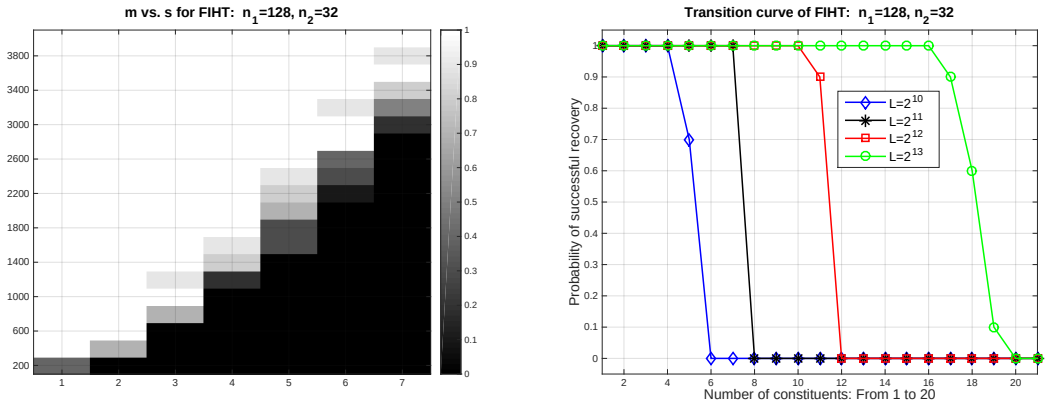

(a) (b)

Figure 4: Empirical phase transitions of FIHT for the application from the IoT.

complex Gaussian variables; and b) $\boldsymbol{C}_k = \boldsymbol{D}_k\boldsymbol{H}$, where $\boldsymbol{H}$ is the first n_2 columns of an $m \times m$ Hadamard matrix, and then premultiplied by an $m \times m$ diagonal random sign matrix $\boldsymbol{D}$. The target signals $\boldsymbol{x}_k$ and $\boldsymbol{h}_k$ are complex Gaussian random vectors. The phase transition plot for $\boldsymbol{C}_k$ being a Gaussian matrix is presented in Figure 4a, where the number of constituents varies from 1 to 7. Figure 4a displays a linear correlation between the number of necessary measurements and the number of the constituents, and it shows that $m \approx 3s(n_1 + n_2)$ number of measurements are sufficient for successful demixing of the transmission signals and the impulse responses. When $\boldsymbol{C}_k$ is formed from a partial Hadamard matrix, we only test $m = 2^{10}, 2^{11}, 2^{12}$, and 2^{13} due to the Hadamard conjecture [10], but a larger range of s are tested. The phase transition plot presented in Figure 4b shows that FIHT is still able to successfully demix the transmission signals and the impulse responses when $m \approx 3s(n_1+n_2)$.

3.4 Robustness to Additive Noise

We explore the robustness of FIHT against additive noise by conducting tests on the problem instances set up in the last three subsections. We have established in Theorem 2.3 that Algorithm 1 (IHT) is stable in the presence of additive noise, but could not provide a similar theoretical robustness guarantee for FIHT because of the more involved nature of the current convergence analysis framework of FIHT. Yet, empirical evidence clearly indicates that FIHT also shows the desirable stability vis-à-vis noise, as we will demonstrate in this subsection.

We assume that the measurement vectors $\boldsymbol{y}$ in the tests are corrupted by

$$\boldsymbol{e} = \sigma \cdot \|\boldsymbol{y}\| \cdot \frac{\boldsymbol{w}}{\|\boldsymbol{w}\|},$$

where $\boldsymbol{w}$ is an $m \times 1$ standard Gaussian vector, either real or complex up to the testing environment, and σ takes nine different values of σ from 10^{-4} to 1. For each problem instance, two values of the number of measurements are tested. The average relative reconstruction

error out in dB against the signal to noise ratio (SNR) is plotted in Figure 5. The figure shows that the relative reconstruction error scales linearly with the noise levels under different measurement schemes. Moreover, as desired, the relative reconstruction error decreases linearly on a log-log scale as the number of measurements increases.

3.5 A Rank Increasing Heuristic

While in some applications the rank of the target matrices is known a priori, for example in pure states demxing and the IoT application, we may not know the matrix rank exactly in other applications. Here, we suggest a rank increasing heuristic when the rank of the matrices is not known. Starting from rank one, we execute FIHT until the residual is not decreasing significantly; and then we successively increase the rank and repeat this procedure until the residual is sufficiently small. We test this heuristic using problems instances from Subsections 3.1 and 3.2. The relative residual plotted against the number of iterations is presented in Figure 6. The jumps in the error curves correspond to the iteration steps when the algorithm increases the assumed rank of the unknown matrices by 1.

4 Proofs

4.1 Proof of Theorem 2.1

Proof. Let $\{\boldsymbol{Z}_k\}_{k=1}^s$ be a set of fixed matrices with $\text{rank}(\boldsymbol{Z}_k) \leq r$ for all $1 \leq k \leq s$. One can easily see that

$$\sqrt{m}\sum_{k=1}^{s}\mathcal{A}_k(\boldsymbol{Z}_k) = \begin{bmatrix} \sum_{k=1}^{s}\langle \boldsymbol{A}_{k,1}, \boldsymbol{Z}_k\rangle \\ \vdots \\ \sum_{k=1}^{s}\langle \boldsymbol{A}_{k,m}, \boldsymbol{Z}_k\rangle \end{bmatrix}.$$

Since for any $1 \leq p \leq m$, $\boldsymbol{A}_{1,p}, \boldsymbol{A}_{2,p}, \cdots, \boldsymbol{A}_{s,p}$ are independent from each other, one has

$$\sum_{k=1}^{s}\langle \boldsymbol{A}_{k,p}, \boldsymbol{Z}_k\rangle \sim \mathcal{N}(0, \sum_{k=1}^{s}\|\boldsymbol{Z}_k\|_F^2).$$

Therefore,

$$m\left\|\sum_{k=1}^{s}\mathcal{A}_k(\boldsymbol{Z}_k)\right\|^2 = \sum_{p=1}^{m}\left|\sum_{k=1}^{s}\langle \boldsymbol{A}_{k,p}, \boldsymbol{Z}_k\rangle\right|^2 \sim \sum_{p=1}^{m}\left(\sum_{k=1}^{s}\|\boldsymbol{Z}_k\|_F^2\right)\xi_p^2,$$

where $\xi_p^2,\ p = 1, \cdots, m$ are i.i.d χ^2 variables. Consequently,

$$\mathbb{E}\left[m\left\|\sum_{k=1}^{s}\mathcal{A}_k(\boldsymbol{Z}_k)\right\|^2\right] = m\sum_{k=1}^{s}\|\boldsymbol{Z}_k\|_F^2,$$

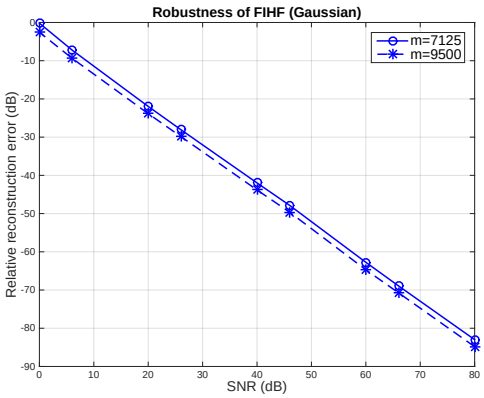

(a)

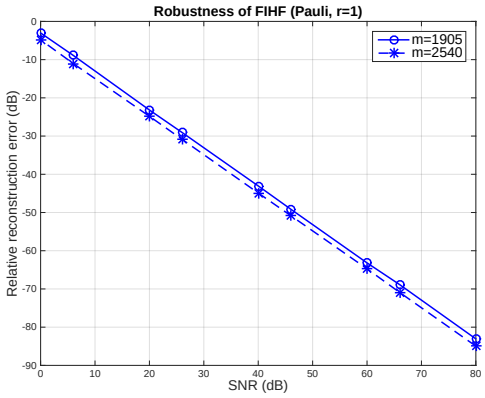

(b)

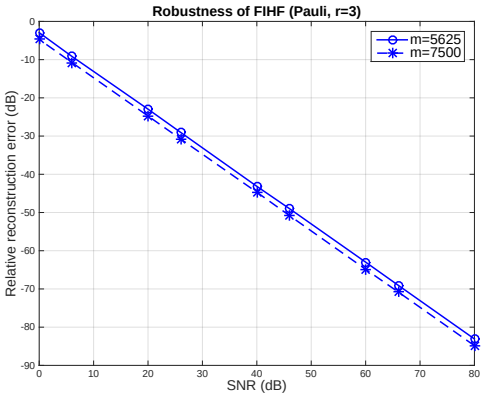

(c)

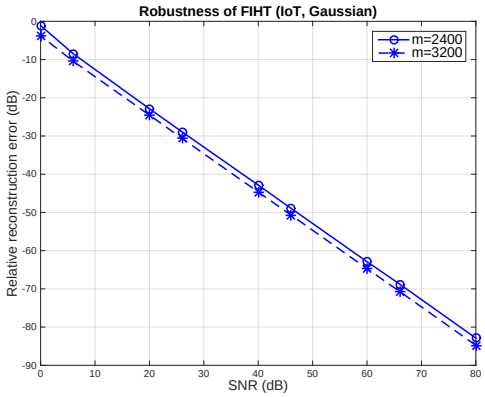

(d)

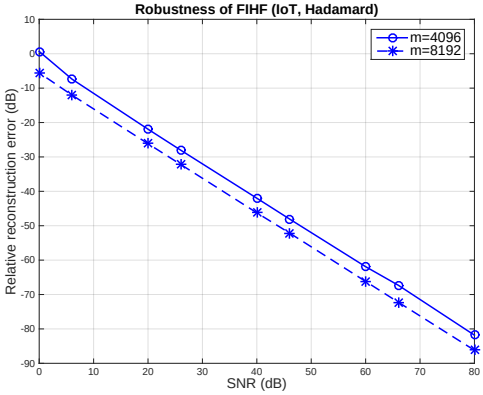

(e)

Figure 5: Robustness of FIHT under different measurements: (a) Gaussian measurements with $s = 5$, (b) and (c) Pauli measurements with $s = 5$, (d) and (e) applications in the IoT with $s = 5$ and $s = 10$, respectively.

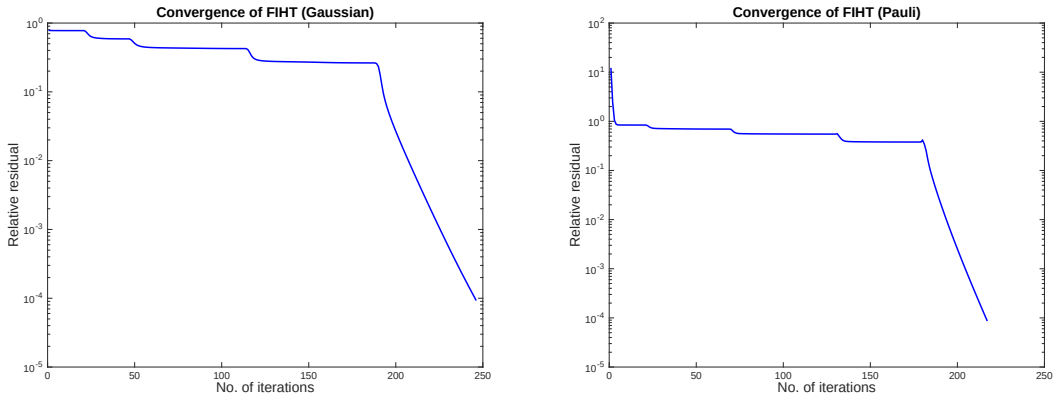

(a) (b)

Figure 6: Rank increasing heuristic for FIHT: (a) Gaussian measurements with $s = 5$, $n = 50$, $r = 5$, and $m = 3sr(2n - r)$, and (b) Pauli measurements with $s = 5$, $n = 2^6$, $r = 5$, and $m = 3sr(2n - r)$. The jumps in the two error curves correspond to the iteration steps when the algorithm increases the assumed rank of the unknown matrices.

and the application of the Bernstein inequality for Chi-squared variables [24] gives

$$\mathbb{P}\left(\left| m\left\|\sum_{k=1}^{s}\mathcal{A}_k(\boldsymbol{Z}_k)\right\|^2 - m\sum_{k=1}^{s}\|\boldsymbol{Z}_k\|_F^2\right| \geq t\right)$$
$$\leq 2\exp\left(-\min\left\{\frac{t^2}{8m\left(\sum_{k=1}^{s}\|\boldsymbol{Z}_k\|_F^2\right)^2}, \frac{t}{8\sum_{k=1}^{s}\|\boldsymbol{Z}_k\|_F^2}\right\}\right).$$

Taking $t = \delta \cdot m\sum_{k=1}^{s}\|\boldsymbol{Z}_k\|_F^2$ gives

$$\mathbb{P}\left(\left|\left\|\sum_{k=1}^{s}\mathcal{A}_k(\boldsymbol{Z}_k)\right\|^2 - \sum_{k=1}^{s}\|\boldsymbol{Z}_k\|_F^2\right| \geq \delta\sum_{k=1}^{s}\|\boldsymbol{Z}_k\|_F^2\right) \leq 2\exp\left(-\frac{m\delta^2}{8}\right)$$

for $0 < \delta < 1$. This implies for fixed $\{\boldsymbol{Z}_k\}_{k=1}^{s}$, we have that

$$(1-\delta)\sum_{k=1}^{s}\|\boldsymbol{Z}_k\|_F^2 \leq \left\|\sum_{k=1}^{s}\mathcal{A}_k(\boldsymbol{Z}_k)\right\|^2 \leq (1+\delta)\sum_{k=1}^{s}\|\boldsymbol{Z}_k\|_F^2 \tag{10}$$

holds with probability at least $1 - 2e^{-m\delta^2/8}$.

The rest of the proof follows essentially from the proof of Theorem 2.3 in [5]. Due to the homogeneity of the problem with respect to $\sum_{k=1}^{s}\|\boldsymbol{Z}_k\|_F^2$, we only need to show (10) for the set of matrices in

$$\mathcal{S}_r = \left\{[\boldsymbol{Z}_1, \cdots, \boldsymbol{Z}_s]^T : \ \text{rank}(\boldsymbol{Z}_k) \leq r \text{ and } \sum_{k=1}^{s}\|\boldsymbol{Z}\|_F^2 = 1\right\}.$$

Based on the SVD of each $\boldsymbol{Z}_k$, we have

$$\mathcal{S}_r = \left\{ [\boldsymbol{U}_1\boldsymbol{\Sigma}_1\boldsymbol{V}_1^*, \cdots, \boldsymbol{U}_s\boldsymbol{\Sigma}_s\boldsymbol{V}_s^*]^T : \ \boldsymbol{U}_k^*\boldsymbol{U}_k = \boldsymbol{I}, \boldsymbol{V}_k^*\boldsymbol{V} = \boldsymbol{I} \text{ and } \sum_{k=1}^{s}\sum_{i=1}^{r} |\boldsymbol{\Sigma}_k^{i,i}|^2 = 1 \right\}.$$

Define $\mathcal{Q}_{n,r} = \{\boldsymbol{Q} \in \mathbb{R}^{n\times r} : \ \boldsymbol{Q}^*\boldsymbol{Q} = \boldsymbol{I}\}$. It has been shown in [5] that $\mathcal{Q}_{n,r}$ has a ϵ-net $\overline{\mathcal{Q}}_{n,r} \subset \mathcal{Q}_{n,r}$ of cardinality $|\overline{\mathcal{Q}}_{n,r}| \leq (3/\epsilon)^{nr}$ under the $\|\cdot\|_{1,2}$ norm. In other words, for any $\boldsymbol{Q} \in \mathcal{Q}_{n,r}$, there exists a $\overline{\boldsymbol{Q}} \in \overline{\mathcal{Q}}_{n,r}$ such that

$$\left\|\boldsymbol{Q} - \overline{\boldsymbol{Q}}\right\|_{1,2} = \max_{1\leq i\leq r} \left\|\boldsymbol{Q}(:,i) - \overline{\boldsymbol{Q}}(:,i)\right\| \leq \epsilon.$$

Let $\mathcal{D}$ be the set of $rs \times rs$ diagonal matrices with unit norm diagonal entries. It is well known that $\mathcal{D}$ has a ϵ-net $\overline{\mathcal{D}} \subset \mathcal{D}$ of cardinality $|\overline{\mathcal{D}}| \leq (3/\epsilon)^{rs}$ [24]. Define

$$\overline{\mathcal{S}}_r = \left\{ [\overline{\boldsymbol{U}}_1\overline{\boldsymbol{\Sigma}}_1\overline{\boldsymbol{V}}_1^*, \cdots, \overline{\boldsymbol{U}}_s\overline{\boldsymbol{\Sigma}}_s\overline{\boldsymbol{V}}_s^*]^T : \ \overline{\boldsymbol{U}}_k \in \overline{\mathcal{Q}}_{n,r}, \ \overline{\boldsymbol{V}}_k \in \overline{\mathcal{Q}}_{n,r}, \text{ and } \operatorname{diag}(\overline{\boldsymbol{\Sigma}}_1, \cdots, \overline{\boldsymbol{\Sigma}}_k) \in \overline{\mathcal{D}} \right\}.$$

We are going to show that $\overline{\mathcal{S}}_r$ is a 3ϵ-net of $\mathcal{S}_r$ with cardinality $|\overline{\mathcal{S}}_r| \leq (3/\epsilon)^{(2n+1)rs}$.

First note that for any $[\boldsymbol{U}_1\boldsymbol{\Sigma}_1\boldsymbol{V}_1^*, \cdots, \boldsymbol{U}_s\boldsymbol{\Sigma}_s\boldsymbol{V}_s^*]^T$, there exists a $[\overline{\boldsymbol{U}}_1\overline{\boldsymbol{\Sigma}}_1\overline{\boldsymbol{V}}_1^*, \cdots, \overline{\boldsymbol{U}}_s\overline{\boldsymbol{\Sigma}}_s\overline{\boldsymbol{V}}_s^*]^T \in \overline{\mathcal{S}}_r$ such that $\left\|\boldsymbol{U}_k - \overline{\boldsymbol{U}}_k\right\|_{1,2} \leq \epsilon$, $\left\|\boldsymbol{V}_k - \overline{\boldsymbol{V}}_k\right\|_{1,2} \leq \epsilon$, and $\sqrt{\sum_{k=1}^{s}\left\|\boldsymbol{\Sigma}_k - \overline{\boldsymbol{\Sigma}}_k\right\|_F^2} \leq \epsilon$. Thus,

$$\begin{aligned}
\left\| \begin{bmatrix} \boldsymbol{U}_1\boldsymbol{\Sigma}_1\boldsymbol{V}_1^* \\ \vdots \\ \boldsymbol{U}_s\boldsymbol{\Sigma}_s\boldsymbol{V}_s^* \end{bmatrix} - \begin{bmatrix} \overline{\boldsymbol{U}}_1\overline{\boldsymbol{\Sigma}}_1\overline{\boldsymbol{V}}_1^* \\ \vdots \\ \overline{\boldsymbol{U}}_s\overline{\boldsymbol{\Sigma}}_s\overline{\boldsymbol{V}}_s^* \end{bmatrix} \right\|_F &\leq \left\| \begin{bmatrix} (\boldsymbol{U}_1 - \overline{\boldsymbol{U}}_1)\boldsymbol{\Sigma}_1\boldsymbol{V}_1^* \\ \vdots \\ (\boldsymbol{U}_s - \overline{\boldsymbol{U}}_s)\boldsymbol{\Sigma}_s\boldsymbol{V}_s^* \end{bmatrix} \right\|_F + \left\| \begin{bmatrix} \overline{\boldsymbol{U}}_1(\boldsymbol{\Sigma}_1 - \overline{\boldsymbol{\Sigma}}_1)\boldsymbol{V}_1^* \\ \vdots \\ \overline{\boldsymbol{U}}_s(\boldsymbol{\Sigma}_s - \overline{\boldsymbol{\Sigma}}_s)\boldsymbol{V}_s^* \end{bmatrix} \right\|_F \\
&\quad + \left\| \begin{bmatrix} \overline{\boldsymbol{U}_1\boldsymbol{\Sigma}}_1(\boldsymbol{V}_1 - \overline{\boldsymbol{V}}_1)^* \\ \vdots \\ \overline{\boldsymbol{U}_s\boldsymbol{\Sigma}}_s(\boldsymbol{V}_s - \overline{\boldsymbol{V}}_s)^* \end{bmatrix} \right\|_F \\
&:= \mathbf{I}_1 + \mathbf{I}_2 + \mathbf{I}_3 \,.
\end{aligned}$$

Since

$$\begin{aligned}
\left\| \begin{bmatrix} (\boldsymbol{U}_1 - \overline{\boldsymbol{U}}_1)\boldsymbol{\Sigma}_1\boldsymbol{V}_1^* \\ \vdots \\ (\boldsymbol{U}_s - \overline{\boldsymbol{U}}_s)\boldsymbol{\Sigma}_s\boldsymbol{V}_s^* \end{bmatrix} \right\|_F^2 &= \sum_{k=1}^{s} \left\|(\boldsymbol{U}_k - \overline{\boldsymbol{U}}_k)\boldsymbol{\Sigma}_k\boldsymbol{V}_k^*\right\|_F^2 \\
&= \sum_{k=1}^{s} \left\|(\boldsymbol{U}_k - \overline{\boldsymbol{U}}_k)\boldsymbol{\Sigma}_k\right\|_F^2 \\
&= \sum_{k=1}^{s}\sum_{i=1}^{r} |\boldsymbol{\Sigma}_k^{i,i}|^2 \left\|\boldsymbol{U}_k(:,i) - \overline{\boldsymbol{U}}_k(:,i)\right\|^2 \\
&\leq \sum_{k=1}^{s}\sum_{i=1}^{r} |\boldsymbol{\Sigma}_k^{i,i}|^2 \left\|\boldsymbol{U}_k - \overline{\boldsymbol{U}}_k\right\|_{1,2}^2
\end{aligned}$$

$$\leq \epsilon^2 \sum_{k=1}^{s} \sum_{i=1}^{r} |\boldsymbol{\Sigma}_k^{i,i}|^2 = \epsilon^2,$$

we have $\mathbf{I}_1 \leq \epsilon$. Similarly, the third term can be bounded as $\mathbf{I}_3 \leq \epsilon$. To bound $\mathbf{I}_2$, note that

$$\left\| \begin{bmatrix} \overline{\boldsymbol{U}}_1(\boldsymbol{\Sigma}_1 - \overline{\boldsymbol{\Sigma}}_1)\boldsymbol{V}_1^* \\ \vdots \\ \overline{\boldsymbol{U}}_s(\boldsymbol{\Sigma}_1 - \overline{\boldsymbol{\Sigma}}_s)\boldsymbol{V}_s^* \end{bmatrix} \right\|_F^2 = \sum_{k=1}^{s} \left\| \overline{\boldsymbol{U}}_k(\boldsymbol{\Sigma}_k - \overline{\boldsymbol{\Sigma}}_k)\boldsymbol{V}_k^* \right\|_F^2$$
$$= \sum_{k=1}^{s} \left\| \boldsymbol{\Sigma}_k - \overline{\boldsymbol{\Sigma}}_k \right\|_F^2 = \epsilon^2.$$

So we also have $\mathbf{I}_2 \leq \epsilon$. Combining the bounds for $\mathbf{I}_1$, $\mathbf{I}_2$ and $\mathbf{I}_3$ together implies that $\overline{\mathcal{S}}_r$ is a 3ϵ-net of $\mathcal{S}_r$ with cardinality $|\overline{\mathcal{S}}_r| \leq (3/\epsilon)^{(2n+1)rs}$. Therefore,

$$\mathbb{P}\left(\max_{\{\boldsymbol{Z}_k\}_{k=1}^s \in \overline{\mathcal{S}}_r} \left| \left\| \sum_{k=1}^{s} \mathcal{A}_k(\boldsymbol{Z}_k) \right\|^2 - \sum_{k=1}^{s} \|\boldsymbol{Z}_k\|_F^2 \right| \geq \delta/2 \right) \leq 2\left(\frac{3}{\epsilon}\right)^{(2n+1)rs} e^{-m\delta^2/32} \leq 2e^{-m\delta^2/64}$$

provided

$$m \geq 64\delta^{-2}(2n+1)rs\log(3/\epsilon).$$

So for all $\left\{\overline{\boldsymbol{Z}}_k\right\} \in \overline{\mathcal{S}}_k$, we have

$$1 - \delta/2 \leq \left\| \sum_{k=1}^{s} \mathcal{A}_k(\overline{\boldsymbol{Z}}_k) \right\| \leq 1 + \delta/2$$

with probability at least $1 - 2e^{-m\delta^2/64}$.

Define

$$\kappa_r = \sup_{\{\boldsymbol{Z}_k\}_{k=1}^s \in \mathcal{S}_r} \left\| \sum_{k=1}^{s} \mathcal{A}_k(\boldsymbol{Z}_k) \right\|.$$

In the following, we will take $\epsilon = \delta/(12\sqrt{2})$. So $\overline{\mathcal{S}}_r$ is a $\delta/(4\sqrt{2})$-net of $\mathcal{S}_r$; and for any $\{\boldsymbol{Z}_k\}_{k=1}^s \in \mathcal{S}_r$, there exist $\{\overline{\boldsymbol{Z}}_k\}_{k=1}^s \in \overline{\mathcal{S}}_r$ such that

$$\left\| \begin{bmatrix} \boldsymbol{Z}_1 - \overline{\boldsymbol{Z}}_1 \\ \vdots \\ \boldsymbol{Z}_s - \overline{\boldsymbol{Z}}_s \end{bmatrix} \right\|_F = \sqrt{\sum_{k=1}^{s} \left\| \boldsymbol{Z}_k - \overline{\boldsymbol{Z}}_k \right\|_F^2} \leq \delta/(4\sqrt{2}).$$

Thus,

$$\left\| \sum_{k=1}^{s} \mathcal{A}_k(\boldsymbol{Z}_k) \right\| \leq \left\| \sum_{k=1}^{s} \mathcal{A}_k\left(\overline{\boldsymbol{Z}}_k\right) \right\| + \left\| \sum_{k=1}^{s} \mathcal{A}_k\left(\boldsymbol{Z}_k - \overline{\boldsymbol{Z}}_k\right) \right\|$$

$$\leq 1 + \delta/2 + \left\| \sum_{k=1}^{s} \mathcal{A}_k \left(\boldsymbol{Z}_k - \overline{\boldsymbol{Z}}_k \right) \right\|. \tag{11}$$

Since $\boldsymbol{Z}_k - \overline{\boldsymbol{Z}}_k$ is a matrix of rank at most $2r$, we can decompose it as

$$\boldsymbol{Z}_k - \overline{\boldsymbol{Z}}_k = \boldsymbol{Y}_k^1 + \boldsymbol{Y}_k^2,$$

where $\langle \boldsymbol{Y}_k^1, \boldsymbol{Y}_k^2 \rangle = 0$, $\operatorname{rank}(\boldsymbol{Y}_k^i) \leq r$ $(i = 1, 2)$, and $\|\boldsymbol{Y}_k^1\|_F^2 + \|\boldsymbol{Y}_k^2\|_F^2 = \left\|\boldsymbol{Z}_k - \overline{\boldsymbol{Z}}_k\right\|_F^2$. Then it follows that

$$\begin{aligned}
\left\| \sum_{k=1}^{s} \mathcal{A}_k \left(\boldsymbol{Z}_k - \overline{\boldsymbol{Z}}_k \right) \right\| &\leq \left\| \sum_{k=1}^{s} \mathcal{A}_k(\boldsymbol{Y}_k^1) \right\| + \left\| \sum_{k=1}^{s} \mathcal{A}_k(\boldsymbol{Y}_k^2) \right\| \\
&\leq \kappa_r \left(\sqrt{\sum_{k=1}^{s} \|\boldsymbol{Y}_k^1\|_F^2} + \sqrt{\sum_{k=1}^{s} \|\boldsymbol{Y}_k^2\|_F^2} \right) \\
&\leq \sqrt{2}\kappa_r \sqrt{\sum_{k=1}^{s} \|\boldsymbol{Y}_k^1\|_F^2 + \sum_{k=1}^{s} \|\boldsymbol{Y}_k^2\|_F^2} \\
&= \sqrt{2}\kappa_r \sqrt{\sum_{k=1}^{s} \left\|\boldsymbol{Z}_k - \overline{\boldsymbol{Z}}_k\right\|_F^2} \\
&\leq \frac{\delta \kappa_r}{4}.
\end{aligned} \tag{12}$$

where the second line follows from $\left\{ \boldsymbol{Y}_k^i / \sqrt{\sum_{k=1}^{s} \|\boldsymbol{Y}_k^i\|_F^2} \right\}_{k=1}^{s} \in \mathcal{S}_r$ for $i = 1, 2$. Inserting (12) into (11) gives

$$\kappa_r \leq 1 + \delta/2 + \delta\kappa_r/4,$$

which implies $\kappa_r \leq 1 + \delta$ and

$$\left\| \sum_{k=1}^{s} \mathcal{A}_k(\boldsymbol{Z}_k) \right\| \leq 1 + \delta \tag{13}$$

for all $\{\boldsymbol{Z}_k\}_{k=1}^{s} \in \mathcal{S}_r$. The lower bound can be obtained in the following way:

$$\begin{aligned}
\left\| \sum_{k=1}^{s} \mathcal{A}_k(\boldsymbol{Z}_k) \right\| &\geq \left\| \sum_{k=1}^{s} \mathcal{A}_k \left(\overline{\boldsymbol{Z}}_k \right) \right\| - \left\| \sum_{k=1}^{s} \mathcal{A}_k \left(\boldsymbol{Z}_k - \overline{\boldsymbol{Z}}_k \right) \right\| \\
&\geq 1 - \delta/2 - \delta(1 + \delta)/4 \\
&\geq 1 - \delta.
\end{aligned} \tag{14}$$

where the second line follows from (12) and the bound for κ_r. The upper and lower bounds in (13) and (14) can be easily transferred into the desired ARIP bounds by a change of variables. □

4.2 Proof of Theorem 2.2

Proof. Let $\boldsymbol{U}_{k,l+1}$, $\boldsymbol{U}_{k,l}$, and $\boldsymbol{U}_k$ be the left singular vectors of $\boldsymbol{X}_{k,l+1}$, $\boldsymbol{X}_{k,l}$ and $\boldsymbol{X}_k$, respectively. Let $\boldsymbol{Q}_{k,l}$ be an $n \times 3r$ orthonormal matrix which spans the union of the column spaces of $\boldsymbol{U}_{k,l+1}$, $\boldsymbol{U}_{k,l}$, and $\boldsymbol{U}_k$. We have $\boldsymbol{X}_{k,l+1} = \mathcal{P}_{\boldsymbol{U}_{k,l+1}}(\boldsymbol{X}_{k,l} + \alpha_{k,l}\boldsymbol{G}_{k,l})$. Moreover, by comparing the following two equalities,

$$\begin{aligned}
\|\boldsymbol{X}_{k,l+1} - (\boldsymbol{X}_{k,l} + \alpha_{k,l}\boldsymbol{G}_{k,l})\|_F^2 &= \left\|\boldsymbol{X}_{k,l+1} - \mathcal{P}_{\boldsymbol{Q}_{k,l}}(\boldsymbol{X}_{k,l} + \alpha_{k,l}\boldsymbol{G}_{k,l})\right\|_F^2 + \left\|\mathcal{P}_{\boldsymbol{Q}_{k,l}^{\perp}}(\boldsymbol{X}_{k,l} + \alpha_{k,l}\boldsymbol{G}_{k,l})\right\|_F^2, \\
\|\boldsymbol{X}_{k} - (\boldsymbol{X}_{k,l} + \alpha_{k,l}\boldsymbol{G}_{k,l})\|_F^2 &= \left\|\boldsymbol{X}_{k} - \mathcal{P}_{\boldsymbol{Q}_{k,l}}(\boldsymbol{X}_{k,l} + \alpha_{k,l}\boldsymbol{G}_{k,l})\right\|_F^2 + \left\|\mathcal{P}_{\boldsymbol{Q}_{k,l}^{\perp}}(\boldsymbol{X}_{k,l} + \alpha_{k,l}\boldsymbol{G}_{k,l})\right\|_F^2,
\end{aligned}$$

we have

$$\left\|\boldsymbol{X}_{k,l+1} - \mathcal{P}_{\boldsymbol{Q}_{k,l}}(\boldsymbol{X}_{k,l} + \alpha_{k,l}\boldsymbol{G}_{k,l})\right\|_F \leq \left\|\boldsymbol{X}_{k} - \mathcal{P}_{\boldsymbol{Q}_{k,l}}(\boldsymbol{X}_{k,l} + \alpha_{k,l}\boldsymbol{G}_{k,l})\right\|_F$$

since $\boldsymbol{X}_{k,l+1}$ is the best rank r approximation of $\boldsymbol{X}_{k,l} + \alpha_{k,l}\boldsymbol{G}_{k,l}$. Thus,

$$\begin{aligned}
\|\boldsymbol{X}_{k,l+1} - \boldsymbol{X}_k\|_F &\leq \left\|\boldsymbol{X}_{k,l+1} - \mathcal{P}_{\boldsymbol{Q}_{k,l}}(\boldsymbol{X}_{k,l} + \alpha_{k,l}\boldsymbol{G}_{k,l})\right\|_F + \left\|\boldsymbol{X}_{k} - \mathcal{P}_{\boldsymbol{Q}_{k,l}}(\boldsymbol{X}_{k,l} + \alpha_{k,l}\boldsymbol{G}_{k,l})\right\|_F \\
&\leq 2\left\|\boldsymbol{X}_{k} - \mathcal{P}_{\boldsymbol{Q}_{k,l}}(\boldsymbol{X}_{k,l} + \alpha_{k,l}\boldsymbol{G}_{k,l})\right\|_F \\
&= 2\left\|\mathcal{P}_{\boldsymbol{Q}_{k,l}}(\boldsymbol{X}_{k,l}) - \mathcal{P}_{\boldsymbol{Q}_{k,l}}(\boldsymbol{X}_{k}) + \alpha_{k,l}\mathcal{P}_{\boldsymbol{Q}_{k,l}}(\boldsymbol{G}_{k,l})\right\|_F.
\end{aligned}$$

Consequently we have

$$\begin{aligned}
\left\|\begin{bmatrix}\boldsymbol{X}_{1,l+1} - \boldsymbol{X}_1 \\ \vdots \\ \boldsymbol{X}_{s,l+1} - \boldsymbol{X}_s\end{bmatrix}\right\|_F &\leq 2\left\|\begin{bmatrix}\mathcal{P}_{\boldsymbol{Q}_{1,l}}(\boldsymbol{X}_{1,l}) - \mathcal{P}_{\boldsymbol{Q}_{1,l}}(\boldsymbol{X}_1) + \alpha_{1,l}\mathcal{P}_{\boldsymbol{Q}_{1,l}}(\boldsymbol{G}_{1,l})\right] \\ \vdots \\ \mathcal{P}_{\boldsymbol{Q}_{s,l}}(\boldsymbol{X}_{s,l}) - \mathcal{P}_{\boldsymbol{Q}_{s,l}}(\boldsymbol{X}_s) + \alpha_{s,l}\mathcal{P}_{\boldsymbol{Q}_{s,l}}(\boldsymbol{G}_{s,l})\end{bmatrix}\right\|_F \\
&= 2\left\|\begin{bmatrix}\mathcal{P}_{\boldsymbol{Q}_{1,l}}(\boldsymbol{X}_{1,l} - \boldsymbol{X}_1) - \alpha_{1,l}\mathcal{P}_{\boldsymbol{Q}_{1,l}}\mathcal{A}_1^*\sum_{k=1}^s \mathcal{A}_k(\boldsymbol{X}_{k,l} - \boldsymbol{X}_k)\right] \\ \vdots \\ \mathcal{P}_{\boldsymbol{Q}_{s,l}}(\boldsymbol{X}_{s,l} - \boldsymbol{X}_s) - \alpha_{s,l}\mathcal{P}_{\boldsymbol{Q}_{1,l}}\mathcal{A}_s^*\sum_{k=1}^s \mathcal{A}_k(\boldsymbol{X}_{k,l} - \boldsymbol{X}_k)\end{bmatrix}\right\|_F \\
&= 2\left\|\begin{bmatrix}\mathcal{P}_{\boldsymbol{Q}_{1,l}}(\boldsymbol{X}_{1,l} - \boldsymbol{X}_1) - \alpha_{1,l}\mathcal{P}_{\boldsymbol{Q}_{1,l}}\mathcal{A}_1^*\sum_{k=1}^s \mathcal{A}_k\mathcal{P}_{\boldsymbol{Q}_{k,l}}(\boldsymbol{X}_{k,l} - \boldsymbol{X}_k)\right] \\ \vdots \\ \mathcal{P}_{\boldsymbol{Q}_{s,l}}(\boldsymbol{X}_{s,l} - \boldsymbol{X}_s) - \alpha_{s,l}\mathcal{P}_{\boldsymbol{Q}_{1,l}}\mathcal{A}_s^*\sum_{k=1}^s \mathcal{A}_k\mathcal{P}_{\boldsymbol{Q}_{k,l}}(\boldsymbol{X}_{k,l} - \boldsymbol{X}_k)\end{bmatrix}\right\|_F \\
&:= 2\|\mathbf{I}_1\|_F
\end{aligned} \tag{15}$$

where in the third line, we have used the fact $\mathcal{P}_{\boldsymbol{Q}_{k,l}}(\boldsymbol{X}_{k,l} - \boldsymbol{X}_k) = (\boldsymbol{X}_{k,l} - \boldsymbol{X}_k)$ for all $1 \leq k \leq s$. In order to bound $\mathbf{I}_1$, we first rewrite it into the following matrix-vector product form

$$\mathbf{I}_1 = \mathcal{A}_\alpha\begin{bmatrix}\boldsymbol{X}_{1,l} - \boldsymbol{X}_1 \\ \vdots \\ \boldsymbol{X}_{s,l} - \boldsymbol{X}_s\end{bmatrix}, \tag{16}$$

where

$$\mathcal{A}_\alpha = \begin{bmatrix} \mathcal{P}_{\boldsymbol{Q}_{1,l}} & \cdots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \cdots & P_{\boldsymbol{Q}_{s,l}} \end{bmatrix} - \begin{bmatrix} \alpha_{1,l} & & \\ & \ddots & \\ & & \alpha_{s,l} \end{bmatrix} \begin{bmatrix} \mathcal{P}_{\boldsymbol{Q}_{1,l}}\mathcal{A}_1^* \\ \vdots \\ \mathcal{P}_{\boldsymbol{Q}_{s,l}}\mathcal{A}_s^* \end{bmatrix} \begin{bmatrix} \mathcal{A}_1\mathcal{P}_{\boldsymbol{Q}_{1,l}} & \cdots & \mathcal{A}_s\mathcal{P}_{\boldsymbol{Q}_{s,l}} \end{bmatrix}.$$

So it suffices to bound the spectral norm of $\mathcal{A}_\alpha$. Denote $\mathcal{A}_\alpha$ by $\mathcal{A}$ when $\alpha_{k,l} = 1$ for all $1 \le k \le s$ in $\mathcal{A}_\alpha$. Let $\boldsymbol{Y} = [\boldsymbol{Y}_1, \cdots, \boldsymbol{Y}_s]^T$. Since $\mathcal{A}$ is self-adjoint, we have

$$\begin{aligned} \|\mathcal{A}\| &= \sup_{\|\boldsymbol{Y}\|_F=1} |\langle \boldsymbol{Y}, \mathcal{A}(\boldsymbol{Y})\rangle| \\ &= \sup_{\|\boldsymbol{Y}\|_F=1} \left| \sum_{k=1}^s \left\|\mathcal{P}_{\boldsymbol{Q}_{k,l}}(\boldsymbol{Y}_k)\right\|_F^2 - \left\| \sum_{k=1}^s \mathcal{A}_k\mathcal{P}_{\boldsymbol{Q}_{k,l}}(\boldsymbol{Y}_k) \right\| \right| \\ &\le \sup_{\|\boldsymbol{Y}\|_F=1} \delta_{3r} \sum_{k=1}^s \left\|\mathcal{P}_{\boldsymbol{Q}_{k,l}}(\boldsymbol{Y}_k)\right\|_F^2 \\ &\le \delta_{3r}, \end{aligned}$$

where the third line follows from the ARIP and the fact $\mathcal{P}_{\boldsymbol{Q}_{k,l}}(\boldsymbol{Y}_k),\ k = 1, \cdots, s$ are matrices of rank at most $3r$. It follows that

$$\begin{aligned} \|\mathcal{A}_\alpha\| &\le \|\mathcal{A}_\alpha - \mathcal{A}\| + \|\mathcal{A}\| \\ &\le \max_k |\alpha_{k,l} - 1| \left\| \begin{bmatrix} \mathcal{P}_{\boldsymbol{Q}_{1,l}}\mathcal{A}_1^* \\ \mathcal{P}_{\boldsymbol{Q}_{2,l}}\mathcal{A}_2^* \\ \vdots \\ \mathcal{P}_{\boldsymbol{Q}_{s,l}}\mathcal{A}_s^* \end{bmatrix} \begin{bmatrix} \mathcal{A}_1\mathcal{P}_{\boldsymbol{Q}_{1,l}} & \mathcal{A}_2\mathcal{P}_{\boldsymbol{Q}_{2,l}} & \cdots & \mathcal{A}_s\mathcal{P}_{\boldsymbol{Q}_{s,l}} \end{bmatrix} \right\| \\ &\le \frac{\delta_{2r}(1+\delta_{3r})}{1-\delta_{2r}} + \delta_{3r} \\ &\le \frac{2\delta_{3r}}{1-\delta_{3r}}, \end{aligned}$$

where in the third line we use the following ARIP bound for $\alpha_{k,l}$

$$\frac{1}{1+\delta_{2r}} \le \alpha_{k,l} = \frac{\|\mathcal{P}_{T_{k,l}}(\boldsymbol{G}_{k,l})\|_F^2}{\|\mathcal{A}_k\mathcal{P}_{T_{k,l}}(\boldsymbol{G}_{k,l})\|_2^2} \le \frac{1}{1-\delta_{2r}}. \tag{17}$$

Combining the spectral of $\mathcal{A}_\alpha$ together with (15) and (16) gives

$$\left\| \begin{bmatrix} \boldsymbol{X}_{1,l+1} - \boldsymbol{X}_1 \\ \vdots \\ \boldsymbol{X}_{s,l+1} - \boldsymbol{X}_s \end{bmatrix} \right\|_F \le \frac{4\delta_{3r}}{1-\delta_{3r}} \left\| \begin{bmatrix} \boldsymbol{X}_{1,l} - \boldsymbol{X}_1 \\ \vdots \\ \boldsymbol{X}_{s,l} - \boldsymbol{X}_s \end{bmatrix} \right\|_F,$$

which completes the proof of the theorem. □

4.3 Proof of Theorem 2.3

Proof. To prove the convergence of IHT under additive noise, we only need to modify (15) slightly as follows:

$$\begin{aligned}
\left\|\begin{bmatrix}\boldsymbol{X}_{1,l+1}-\boldsymbol{X}_1\\ \vdots\\ \boldsymbol{X}_{s,l+1}-\boldsymbol{X}_s\end{bmatrix}\right\|_F &\le 2\left\|\begin{bmatrix}\mathcal{P}_{\boldsymbol{Q}_{1,l}}(\boldsymbol{X}_{1,l})-\mathcal{P}_{\boldsymbol{Q}_{1,l}}(\boldsymbol{X}_1)+\alpha_{1,l}\mathcal{P}_{\boldsymbol{Q}_{1,l}}(\boldsymbol{G}_{1,l})\\ \vdots\\ \mathcal{P}_{\boldsymbol{Q}_{s,l}}(\boldsymbol{X}_{s,l})-\mathcal{P}_{\boldsymbol{Q}_{s,l}}(\boldsymbol{X}_s)+\alpha_{s,l}\mathcal{P}_{\boldsymbol{Q}_{s,l}}(\boldsymbol{G}_{s,l})\end{bmatrix}\right\|_F\\
&\le 2\left\|\begin{bmatrix}\mathcal{P}_{\boldsymbol{Q}_{1,l}}(\boldsymbol{X}_{1,l}-\boldsymbol{X}_1)-\alpha_{1,l}\mathcal{P}_{\boldsymbol{Q}_{1,l}}\mathcal{A}_1^*\sum_{k=1}^s\mathcal{A}_k(\boldsymbol{X}_{k,l}-\boldsymbol{X}_k)\\ \vdots\\ \mathcal{P}_{\boldsymbol{Q}_{s,l}}(\boldsymbol{X}_{s,l}-\boldsymbol{X}_s)-\alpha_{s,l}\mathcal{P}_{\boldsymbol{Q}_{1,l}}\mathcal{A}_s^*\sum_{k=1}^s\mathcal{A}_k(\boldsymbol{X}_{k,l}-\boldsymbol{X}_k)\end{bmatrix}\right\|_F\\
&\quad+2\left\|\begin{bmatrix}\alpha_{1,l}\mathcal{P}_{\boldsymbol{Q}_{1,l}}\mathcal{A}_1^*(\boldsymbol{e})\\ \vdots\\ \alpha_{s,l}\mathcal{P}_{\boldsymbol{Q}_{s,l}}\mathcal{A}_s^*(\boldsymbol{e})\end{bmatrix}\right\|_F\\
&:=\mathbf{I}'_{11}+\mathbf{I}'_{12}\,.
\end{aligned}$$

The proof of Theorem 2.2 shows that

$$\mathbf{I}'_{11}\le\frac{4\delta_{3r}}{1-\delta_{3r}}\left\|\begin{bmatrix}\boldsymbol{X}_{1,l}-\boldsymbol{X}_1\\ \vdots\\ \boldsymbol{X}_{s,l}-\boldsymbol{X}_s\end{bmatrix}\right\|_F,$$

while $\mathbf{I}'_{12}$ can be bounded as

$$\begin{aligned}
\mathbf{I}'_{12}&\le 2\max_k|\alpha_{k,l}|\left\|\begin{bmatrix}\mathcal{P}_{\boldsymbol{Q}_{1,l}}\mathcal{A}_1^*(\boldsymbol{e})\\ \vdots\\ \mathcal{P}_{\boldsymbol{Q}_{s,l}}\mathcal{A}_s^*(\boldsymbol{e})\end{bmatrix}\right\|_F\\
&=2\max_k|\alpha_{k,l}|\max_{\{\boldsymbol{Y}_k\}_{k=1}^s,\sum_{k=1}^s\|\boldsymbol{Y}_k\|_F^2=1}\sum_{k=1}^s\left\langle\mathcal{P}_{\boldsymbol{Q}_{k,l}}\mathcal{A}_1^*(\boldsymbol{e}),\boldsymbol{Y}_k\right\rangle\\
&\le\frac{2}{1-\delta_{2r}}\|\boldsymbol{e}\|\max_{\{\boldsymbol{Y}_k\}_{k=1}^s,\sum_{k=1}^s\|\boldsymbol{Y}_k\|_F^2=1}\left\|\sum_{k=1}^s\mathcal{A}_k\mathcal{P}_{Q_{k,l}}(\boldsymbol{Y}_k)\right\|\\
&\le\frac{2\sqrt{1+\delta_{3r}}}{1-\delta_{2r}}\|\boldsymbol{e}\|,
\end{aligned}$$

where in the third line we utilize the bound for $\alpha_{k,l}$ in (17) and in the last line we utilize the definition of the ARIP. Therefore,

$$\left\|\begin{bmatrix}\boldsymbol{X}_{1,l+1}-\boldsymbol{X}_1\\ \vdots\\ \boldsymbol{X}_{s,l+1}-\boldsymbol{X}_s\end{bmatrix}\right\|_F\le\frac{4\delta_{3r}}{1-\delta_{3r}}\left\|\begin{bmatrix}\boldsymbol{X}_{1,l}-\boldsymbol{X}_1\\ \vdots\\ \boldsymbol{X}_{s,l}-\boldsymbol{X}_s\end{bmatrix}\right\|_F+\frac{2\sqrt{1+\delta_{3r}}}{1-\delta_{2r}}\|\boldsymbol{e}\|.$$

The proof is complete after we apply the above inequality recursively. □

4.4 Proof of Theorem 2.4

Proof. Let $\boldsymbol{W}_{k,l} = \boldsymbol{X}_{k,l} + \alpha_{k,l}\mathcal{P}_{T_{k,l}}(\boldsymbol{G}_{k,l})$. The following inequality holds:

$$\left\|\begin{bmatrix}\boldsymbol{X}_{1,l+1}-\boldsymbol{X}_1\\ \vdots\\ \boldsymbol{X}_{s,l+1}-\boldsymbol{X}_s\end{bmatrix}\right\|_F = \sqrt{\sum_{k=1}^{s}\|\boldsymbol{X}_{k,l+1}-\boldsymbol{X}_k\|_F^2} \le \sqrt{4\sum_{k=1}^{s}\|\boldsymbol{W}_{k,l}-\boldsymbol{X}_k\|_F^2} = 2\left\|\begin{bmatrix}\boldsymbol{W}_{1,l}-\boldsymbol{X}_1\\ \vdots\\ \boldsymbol{W}_{s,l}-\boldsymbol{X}_s\end{bmatrix}\right\|_F,$$

where the inequality follows from

$$\|\boldsymbol{X}_{k,l+1}-\boldsymbol{X}_k\|_F \le \|\boldsymbol{X}_{k,l+1}-\boldsymbol{W}_{k,l}\|_F + \|\boldsymbol{W}_{k,l}-\boldsymbol{X}_k\|_F \le 2\|\boldsymbol{W}_{k,l}-\boldsymbol{X}_k\|_F.$$

This leads to

$$\begin{aligned}
\left\|\begin{bmatrix}\boldsymbol{X}_{1,l+1}-\boldsymbol{X}_1\\ \vdots\\ \boldsymbol{X}_{s,l+1}-\boldsymbol{X}_s\end{bmatrix}\right\|_F &\le 2\left\|\begin{bmatrix}\boldsymbol{X}_{1,l}+\alpha_{1,l}\mathcal{P}_{T_{1,l}}(\boldsymbol{G}_{1,l})-\boldsymbol{X}_1\\ \vdots\\ \boldsymbol{X}_{s,l}+\alpha_{s,l}\mathcal{P}_{T_{1,l}}(\boldsymbol{G}_{s,l})-\boldsymbol{X}_s\end{bmatrix}\right\|_F\\
&= 2\left\|\begin{bmatrix}\boldsymbol{X}_{1,l}-\boldsymbol{X}_1-\alpha_{1,l}\mathcal{P}_{T_{1,l}}\mathcal{A}_1^*\sum_{k=1}^{s}\mathcal{A}_k(\boldsymbol{X}_{k,l}-\boldsymbol{X}_k)\\ \vdots\\ \boldsymbol{X}_{s,l}-\boldsymbol{X}_s-\alpha_{s,l}\mathcal{P}_{T_{s,l}}\mathcal{A}_s^*\sum_{k=1}^{s}\mathcal{A}_k(\boldsymbol{X}_{k,l}-\boldsymbol{X}_k)\end{bmatrix}\right\|_F\\
&\le 2\left\|\begin{bmatrix}\mathcal{P}_{T_{1,l}}(\boldsymbol{X}_{1,l}-\boldsymbol{X}_1)-\alpha_{1,l}\mathcal{P}_{T_{1,l}}\mathcal{A}_1^*\sum_{k=1}^{s}\mathcal{A}_k\mathcal{P}_{T_{k,l}}(\boldsymbol{X}_{k,l}-\boldsymbol{X}_k)\\ \vdots\\ \mathcal{P}_{T_{s,l}}(\boldsymbol{X}_{s,l}-\boldsymbol{X}_s)-\alpha_{s,l}\mathcal{P}_{T_{s,l}}\mathcal{A}_s^*\sum_{k=1}^{s}\mathcal{A}_k\mathcal{P}_{T_{k,l}}(\boldsymbol{X}_{k,l}-\boldsymbol{X}_k)\end{bmatrix}\right\|_F\\
&\quad + 2\left\|\begin{bmatrix}\alpha_{1,l}\mathcal{P}_{T_{1,l}}\mathcal{A}_1^*\sum_{k=1}^{s}\mathcal{A}_k(\mathcal{I}-\mathcal{P}_{T_{k,l}})(\boldsymbol{X}_{k,l}-\boldsymbol{X}_k)\\ \vdots\\ \alpha_{s,l}\mathcal{P}_{T_{s,l}}\mathcal{A}_s^*\sum_{k=1}^{s}\mathcal{A}_k(\mathcal{I}-\mathcal{P}_{T_{k,l}})(\boldsymbol{X}_{k,l}-\boldsymbol{X}_k)\end{bmatrix}\right\|_F\\
&\quad + 2\left\|\begin{bmatrix}(\mathcal{I}-\mathcal{P}_{T_{1,l}})\boldsymbol{X}_1\\ \vdots\\ (\mathcal{I}-\mathcal{P}_{T_{s,l}})\boldsymbol{X}_s\end{bmatrix}\right\|_F\\
&:= 2\|\mathbf{I}_2\|_F + 2\|\mathbf{I}_3\|_F + 2\|\mathbf{I}_4\|_F. \qquad (18)
\end{aligned}$$

Following the same argument for the bound of $\mathbf{I}_1$ in (15), we can bound $\mathbf{I}_2$ as

$$\|\mathbf{I}_2\|_F \le \frac{2\delta_{2r}}{1-\delta_{2r}}\left\|\begin{bmatrix}\boldsymbol{X}_{1,l}-\boldsymbol{X}_1\\ \vdots\\ \boldsymbol{X}_{s,l}-\boldsymbol{X}_s\end{bmatrix}\right\|_F$$

by noting that all the matrices in $T_{k,l}$ are of rank at most $2r$.

Next, $\mathbf{I}_3$ can be bounded in the following way:

$$\|\mathbf{I}_3\|_F \le \max_k|\alpha_{k,l}| \sup_{\{\boldsymbol{Y}_k\}_{k=1}^s,\sum_{k=1}^s\|\boldsymbol{Y}_k\|_F^2=1}\left|\sum_{k=1}^{s}\left\langle \boldsymbol{Y}_k, \mathcal{P}_{T_{k,l}}\mathcal{A}_k^*\sum_{k=1}^{s}\mathcal{A}_k(\mathcal{I}-\mathcal{P}_{T_{k,l}})(\boldsymbol{X}_{k,l}-\boldsymbol{X}_k)\right\rangle\right|$$

$$
\begin{aligned}
&= \max_k |\alpha_{k,l}| \sup_{\{\boldsymbol{Y}_k\}_{k=1}^s, \sum_{k=1}^s \|\boldsymbol{Y}_k\|_F^2 = 1} \left| \sum_{k=1}^s \left\langle \mathcal{A}_k \mathcal{P}_{T_{k,l}}(\boldsymbol{Y}_k), \sum_{k=1}^s \mathcal{A}_k (\mathcal{I} - \mathcal{P}_{T_{k,l}})(\boldsymbol{X}_{k,l} - \boldsymbol{X}_k) \right\rangle \right| \\
&\leq \max_k |\alpha_{k,l}| \sup_{\{\boldsymbol{Y}_k\}_{k=1}^s, \sum_{k=1}^s \|\boldsymbol{Y}_k\|_F^2 = 1} \delta_{3r} \sqrt{\sum_{k=1}^s \left\| \mathcal{P}_{T_{k,l}}(\boldsymbol{Y}_k) \right\|_F^2} \sqrt{\sum_{k=1}^s \left\| (\mathcal{I} - \mathcal{P}_{T_{k,l}})(\boldsymbol{X}_{k,l} - \boldsymbol{X}_k) \right\|_F^2} \\
&\leq \frac{\delta_{3r}}{1 - \delta_{2r}} \left\| \begin{bmatrix} \boldsymbol{X}_{1,l} - \boldsymbol{X}_1 \\ \vdots \\ \boldsymbol{X}_{s,l} - \boldsymbol{X}_s \end{bmatrix} \right\|_F,
\end{aligned}
$$

where the third line follows from Lemma 5.1 in the appendix, and the fourth line follows from the ARIP bound for $\alpha_{k,l}$ in (17).

To bound $\mathbf{I}_4$, first note that the application of Lemma 5.2 in the appendix gives

$$
\|(\mathcal{I} - \mathcal{P}_{T_{k,l}})(\boldsymbol{X}_k)\|_F \leq \frac{\|\boldsymbol{X}_{k,l} - \boldsymbol{X}_k\|_F^2}{\sigma_{\min}(\boldsymbol{X}_k)} \leq \frac{\max_k \|\boldsymbol{X}_{k,l} - \boldsymbol{X}_k\|_F}{\sigma_{\min}} \|\boldsymbol{X}_{k,l} - \boldsymbol{X}_k\|_F,
$$

where $\sigma_{\min} := \min_k \sigma_{\min}(\boldsymbol{X}_k)$. Thus,

$$
\|\mathbf{I}_4\|_F \leq \frac{\max_k \|\boldsymbol{X}_{k,l} - \boldsymbol{X}_k\|_F}{\sigma_{\min}} \left\| \begin{bmatrix} \boldsymbol{X}_{1,l} - \boldsymbol{X}_1 \\ \vdots \\ \boldsymbol{X}_{s,l} - \boldsymbol{X}_s \end{bmatrix} \right\|_F \leq \frac{1}{\sigma_{\min}} \left\| \begin{bmatrix} \boldsymbol{X}_{1,l} - \boldsymbol{X}_1 \\ \vdots \\ \boldsymbol{X}_{s,l} - \boldsymbol{X}_s \end{bmatrix} \right\|_F \cdot \left\| \begin{bmatrix} \boldsymbol{X}_{1,l} - \boldsymbol{X}_1 \\ \vdots \\ \boldsymbol{X}_{s,l} - \boldsymbol{X}_s \end{bmatrix} \right\|_F.
$$

Substituting the bounds for $\mathbf{I}_2$, $\mathbf{I}_3$ and $\mathbf{I}_4$ into (18) gives

$$
\left\| \begin{bmatrix} \boldsymbol{X}_{1,l+1} - \boldsymbol{X}_1 \\ \vdots \\ \boldsymbol{X}_{s,l+1} - \boldsymbol{X}_s \end{bmatrix} \right\|_F \leq 2 \left(\frac{2\delta_{2r}}{1 - \delta_{2r}} + \frac{\delta_{3r}}{1 - \delta_{2r}} + \frac{1}{\sigma_{\min}} \left\| \begin{bmatrix} \boldsymbol{X}_{1,l} - \boldsymbol{X}_1 \\ \vdots \\ \boldsymbol{X}_{s,l} - \boldsymbol{X}_s \end{bmatrix} \right\|_F \right) \left\| \begin{bmatrix} \boldsymbol{X}_{1,l} - \boldsymbol{X}_1 \\ \vdots \\ \boldsymbol{X}_{s,l} - \boldsymbol{X}_s \end{bmatrix} \right\|_F.
$$

Assume

$$
\rho := 2 \left(\frac{2\delta_{2r}}{1 - \delta_{2r}} + \frac{\delta_{3r}}{1 - \delta_{2r}} + \frac{1}{\sigma_{\min}} \left\| \begin{bmatrix} \boldsymbol{X}_{1,0} - \boldsymbol{X}_1 \\ \vdots \\ \boldsymbol{X}_{s,0} - \boldsymbol{X}_s \end{bmatrix} \right\|_F \right) < 1.
$$

By induction, we have

$$
\left\| \begin{bmatrix} \boldsymbol{X}_{1,l+1} - \boldsymbol{X}_1 \\ \vdots \\ \boldsymbol{X}_{s,l+1} - \boldsymbol{X}_s \end{bmatrix} \right\|_F \leq \rho \left\| \begin{bmatrix} \boldsymbol{X}_{1,l} - \boldsymbol{X}_1 \\ \vdots \\ \boldsymbol{X}_{s,l} - \boldsymbol{X}_s \end{bmatrix} \right\|_F
$$

Note that the initial guess is obtained by one-step hard thresholding with $\alpha_{k,l} = 1$. So following the proof of Theorem 2.2 we have

$$
\left\| \begin{bmatrix} \boldsymbol{X}_{1,0} - \boldsymbol{X}_1 \\ \vdots \\ \boldsymbol{X}_{s,0} - \boldsymbol{X}_s \end{bmatrix} \right\|_F \leq 2\delta_{3r} \left\| \begin{bmatrix} \boldsymbol{X}_1 \\ \vdots \\ \boldsymbol{X}_s \end{bmatrix} \right\|_F.
$$

Therefore, one has $\rho < 1$ if

$$2\left(\frac{2\delta_{2r}}{1-\delta_{2r}} + \frac{\delta_{3r}}{1-\delta_{2r}} + \frac{2\delta_{3r}}{\sigma_{\min}}\left\|\begin{bmatrix}\boldsymbol{X}_1\\ \vdots \\ \boldsymbol{X}_s\end{bmatrix}\right\|_F\right) < 1. \tag{19}$$

Since

$$\left\|\begin{bmatrix}\boldsymbol{X}_1\\ \vdots \\ \boldsymbol{X}_s\end{bmatrix}\right\|_F = \sqrt{\sum_{k=1}^{s}\|\boldsymbol{X}_k\|_F^2} \le \sqrt{\sum_{k=1}^{s} r\,{\sigma_{\max}}^2(\boldsymbol{X}_k)} \le \sqrt{rs}\,\sigma_{\max},$$

(19) holds if

$$2\left(\frac{2\delta_{2r}}{1-\delta_{2r}} + \frac{\delta_{3r}}{1-\delta_{2r}} + 2\delta_{3r}\sqrt{rs}\frac{\sigma_{\max}}{\sigma_{\min}}\right) < 1,$$

where $\sigma_{\max} := \max_k \sigma_{\max}(\boldsymbol{X}_k)$. □

5 Conclusion and Future Directions

We have presented the first computationally efficient algorithms that can extract low rank matrices from a sum of linear measurements. These algorithms have potential applications in areas such as wireless communications and quantum tomography. Numerical simulations show an empirical performance that is quite close to the information theoretic limit in terms of demixing ability with respect to the number of measurements.

At the same time, there are still a number of open questions for further directions. Firstly, the robustness analysis of FIHT needs to be addressed in the future. Secondly, our theoretical framework so far is still a bit restrictive, since it only yields close-to-optimal results for Gaussian measurement matrices, which are however of limited use in applications. Thus, one future challenge consist in establishing good Amalgam-RIP bounds for structured measurement matrices. Thirdly, it is also interesting to see whether the Amalgam-RIP can be used to analyse other approaches for low rank matrix demixing. In particular, we want to investigate whether the Amalgam form of RIP in (5), but restricted onto a local subspace, is sufficient for the guarantee analysis of the nuclear norm minimization studied in [14]. Finally, in this paper the Amalgam-RIP is defined for homogeneous data, i.e., matrices of low rank. It is likely that similar Amalgam-RIP can be established for heterogeneous data and then be used in the analysis of different reconstruction programs.

Appendix

Lemma 5.1. *Suppose* $\langle \boldsymbol{Y}_k, \boldsymbol{Z}_k\rangle = 0$ *and* $\operatorname{rank}(\boldsymbol{Y}_k + \boldsymbol{Z}_k) \le c$ *for all* $1 \le k \le s$. *Then*

$$\left|\left\langle \sum_{k=1}^{s}\mathcal{A}_k(\boldsymbol{Y}_k), \sum_{k=1}^{s}\mathcal{A}_k(\boldsymbol{Z}_k)\right\rangle\right| \le \delta_c \sqrt{\sum_{k=1}^{s}\|\boldsymbol{Y}_k\|_F^2}\sqrt{\sum_{k=1}^{s}\|\boldsymbol{Z}_k\|_F^2}.$$

Proof. Due to the homogeneity, we can assume $\sqrt{\sum_{k=1}^{s} \|\boldsymbol{Y}_k\|_F^2} = 1$, $\sqrt{\sum_{k=1}^{s} \|\boldsymbol{Z}_k\|_F^2} = 1$. Since $\langle \boldsymbol{Y}_k, \boldsymbol{Z}_k \rangle = 0$, we have $\sum_{k=1}^{s} \|\boldsymbol{Y}_k + \boldsymbol{Z}_k\|_F^2 = 2$. So

$$2(1-\delta_c) \leq \left\| \sum_{k=1}^{s} \mathcal{A}_k(\boldsymbol{Y}_k \pm \boldsymbol{Z}_k) \right\|^2 \leq 2(1+\delta_c)$$

following from Def. 2.1. Then the application of the parallelogram identity implies

$$\left| \left\langle \sum_{k=1}^{s} \mathcal{A}_k(\boldsymbol{Y}_k), \sum_{k=1}^{s} \mathcal{A}_k(\boldsymbol{Z}_k) \right\rangle \right| = \frac{1}{4} \left| \left\| \sum_{k=1}^{s} \mathcal{A}_k(\boldsymbol{Y}_k + \boldsymbol{Z}_k) \right\|^2 - \left\| \sum_{k=1}^{s} \mathcal{A}_k(\boldsymbol{Y}_k - \boldsymbol{Z}_k) \right\|^2 \right| \leq \delta_c,$$

which concludes the proof. □

Lemma 5.2 ([27, Lem. 4.1])**.** *Let $\boldsymbol{X}_{k,l} = \boldsymbol{U}_{k,l}\boldsymbol{\Sigma}_{k,l}\boldsymbol{V}_{k,l}^*$ be rank r matrix and $T_{k,l}$ be the tangent space of the rank r matrix manifold at $\boldsymbol{X}_{k,l}$. Let $\boldsymbol{X}_k$ be another rank r matrix. One has*

$$\|(\mathcal{I} - \mathcal{P}_{T_{k,l}})(\boldsymbol{X}_k)\|_F \leq \frac{\|\boldsymbol{X}_{k,l} - \boldsymbol{X}_k\|_F^2}{\sigma_{\min}(\boldsymbol{X}_k)}.$$

References

[1] P.-A. Absil, R. Mahony, and R. Sepulchre, *Optimization Algorithms on Matrix Manifolds*, Princeton University Press, 2008.

[2] J.-F. Cai, E. J. Candès, and Z. Shen, *A singular value thresholding algorithm for matrix completion*, SIAM Journal on Optimization, 20 (2010), pp. 1956–1982.

[3] P. Campisi and K. Egiazarian, *Blind image deconvolution: theory and applications*, CRC Press, 2007.

[4] E. J. Candès, X. Li, Y. Ma, and J. Wright, *Robust principal component analysis?*, Journal of the ACM, 58 (2011), pp. 1–37.

[5] E. J. Candès and Y. Plan, *Tight oracle bounds for low-rank matrix recovery from a minimal number of random measurements*, IEEE Transactions on Information Theory, 57 (2011), pp. 2342–2359.

[6] E. J. Candès and B. Recht, *Exact matrix completion via convex optimization*, Foundations of Computational Mathematics, 9 (2009), pp. 717–772.

[7] Y. Chen and X. Ye, *Projection onto a simplex*, arXiv preprint arXiv:1101.6081, (2011).

[8] A. Deville and Y. Deville, *Concepts and criteria for blind quantum source separation*, arXiv preprint arXiv:1611.04002, (2016).

[9] A. Forsgren, P. E. Gill, and M. H. Wright, *Interior methods for nonlinear optimization*, SIAM Review, 44 (2002), pp. 525–597.

[10] A. Hedayat and W. D. Wallis, *Hadamard matrices and their applications*, Annals of Statistics, 6 (1978), pp. 1184–1238.

[11] N. J. Higham, *Computing a nearest symmetric positive semidefinite matrix*, Linear algebra and its applications, 103 (1988), pp. 103–118.

[12] F. Huang and A. Anandkumar, *Convolutional dictionary learning through tensor factorization*, arXiv preprint, arXiv:1506.03509, (2015).

[13] M. Kliesch, R. Kueng, J. Eisert, and D. Gross, *Guaranteed recovery of quantum processes from few measurements*, arXiv preprint arXiv:1701.03135, (2017).

[14] S. Ling and T. Strohmer, *Blind deconvolution meets blind demixing: Algorithms and performance bounds*, IEEE Transactions on Information Theory, (2016, under review).

[15] ——, *Regularized gradient descent: A nonconvex recipe for fast joint blind deconvolution and demixing*, Preprint, [arxiv:1703.08642], (2017).

[16] J. Liu, J. Xin, Y. Qi, F.-G. Zheng, et al., *A time domain algorithm for blind separation of convolutive sound mixtures and L_1 constrained minimization of cross correlations*, Communications in Mathematical Sciences, 7 (2009), pp. 109–128.

[17] M. B. McCoy, V. Cevher, Q. T. Dinh, A. Asaei, and L. Baldassarre, *Convexity in source separation: Models, geometry, and algorithms*, Signal Processing Magazine, IEEE, 31 (2014), pp. 87–95.

[18] M. B. McCoy and J. A. Tropp, *Achievable performance of convex demixing*, tech. rep., Caltech, 2017, Paper dated Feb. 2013. ACM Technical Report 2017-02.

[19] B. Recht, M. Fazel, and P. A. Parrilo, *Guaranteed minimum-rank solutions of linear matrix equations via nuclear norm minimization*, SIAM Review, 52 (2010), pp. 471–501.

[20] J.-L. Starck, M. Elad, and D. L. Donoho, *Image decomposition via the combination of sparse representations and a variational approach*, IEEE transactions on image processing, 14 (2005), pp. 1570–1582.

[21] D. Stöger, P. Jung, and F. Krahmer, *Blind deconvolution and compressed sensing*, in Compressed Sensing Theory and its Applications to Radar, Sonar and Remote Sensing (CoSeRa), 2016 4th International Workshop on, IEEE, 2016, pp. 24–27.

[22] J. Tanner and K. Wei, *Normalized iterative hard thresholding for matrix completion*, SIAM Journal on Scientific Computing, 35 (2013), pp. S104–S125.

[23] B. Vandereycken, *Low rank matrix completion by Riemannian optimization*, SIAM Journal on Optimization, 23 (2013), pp. 1214–1236.

[24] R. Vershynin, *Introduction to the non-asymptotic analysis of random matrices*, in Compressed Sensing: Theory and Applications, Y. C. Eldar and G. Kutyniok, eds., Cambridge University Press, 2012.

[25] X. Wang and H. V. Poor, *Blind equalization and multiuser detection in dispersive CDMA channels*, Communications, IEEE Transactions on, 46 (1998), pp. 91–103.

[26] K. Wei, J.-F. Cai, T. F. Chan, and S. Leung, *Guarantees of Riemannian optimization for low rank matrix completion.* arXiv:1603.06610, 2015.

[27] ——, *Guarantees of Riemannian optimization for low rank matrix recovery*, SIAM J. on Matrix Analysis and Applications, 37 (2016), pp. 1198–1222.

[28] J. Wright, A. Ganesh, K. Min, and Y. Ma, *Compressive principal component pursuit*, Information and Inference, 2 (2013), pp. 32–68.

[29] G. Wunder, H. Boche, T. Strohmer, and P. Jung, *Sparse signal processing concepts for efficient 5G system design*, IEEE Access, 3 (2015), pp. 195–208.

[30] X. Zhou, C. Yang, H. Zhao, and W. Yu, *Low-rank modeling and its applications in image analysis*, ACM Computing Surveys (CSUR), 47 (2015), p. 36.